

AD-A252 548



2



TECHNICAL REPORT RD-GC-92-32

AD-E 951 839

REAL TIME EXECUTIVE FOR MISSILE SYSTEMS
USER'S GUIDE 180386 C INTERFACE

Wanda M. Hughes and
Phillip R. Acuff
Guidance and Control Directorate
Research, Development, and Engineering Center

and

On-line Applications Research Corporation
3315 Memorial Parkway SW
Huntsville, AL 35801

MAY 1992

DTIC
ELECTE
JUN 24 1992
S D



U.S. ARMY MISSILE COMMAND

Redstone Arsenal, Alabama 35898-5000

Approved for Public Release.

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

92 6 2 0 0 0

92-16547



DESTRUCTION NOTICE

FOR CLASSIFIED DOCUMENTS, FOLLOW THE PROCEDURES IN DoD 5200.22-M, INDUSTRIAL SECURITY MANUAL, SECTION II-19 OR DoD 5200.1-R, INFORMATION SECURITY PROGRAM REGULATION, CHAPTER IX. FOR UNCLASSIFIED, LIMITED DOCUMENTS, DESTROY BY ANY METHOD THAT WILL PREVENT DISCLOSURE OF CONTENTS OR RECONSTRUCTION OF THE DOCUMENT.

DISCLAIMER


THE FINDINGS IN THIS REPORT ARE NOT TO BE CONSTRUED AS AN OFFICIAL DEPARTMENT OF THE ARMY POSITION UNLESS SO DESIGNATED BY OTHER AUTHORIZED DOCUMENTS.

TRADE NAMES

USE OF TRADE NAMES OR MANUFACTURERS IN THIS REPORT DOES NOT CONSTITUTE AN OFFICIAL ENDORSEMENT OR APPROVAL OF THE USE OF SUCH COMMERCIAL HARDWARE OR SOFTWARE.

Real Time Executive for Missile Systems

User's Guide i80386 C Interface



U.S. ARMY MISSILE COMMAND
Redstone Arsenal, Alabama 35898-5254

Release 1.31
December 1991

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
1a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED			1b. RESTRICTIVE MARKINGS		
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION / AVAILABILITY OF REPORT Approved for Public Release		
2b. DECLASSIFICATION / DOWNGRADING SCHEDULE					
4. PERFORMING ORGANIZATION REPORT NUMBER(S) Technical Report RD-GC-92-32			5. MONITORING ORGANIZATION REPORT NUMBER(S)		
6a. NAME OF PERFORMING ORGANIZATION Guidance & Control Directorate RD&E Center		6b. OFFICE SYMBOL (if applicable)	7a. NAME OF MONITORING ORGANIZATION		
6c. ADDRESS (City, State, and ZIP Code) Commander, U.S. Army Missile Command ATTN: AMSMI-RD-GC-S Redstone Arsenal, AL 35898-5254			7b. ADDRESS (City, State, and ZIP Code)		
8a. NAME OF FUNDING / SPONSORING ORGANIZATION		8b. OFFICE SYMBOL (if applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER		
8c. ADDRESS (City, State, and ZIP Code)			10. SOURCE OF FUNDING NUMBERS		
			PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO.
11. TITLE (Include Security Classification) Real Time Executive For Missile Systems User's Guide i80386 C Interface					
12. PERSONAL AUTHOR(S) Wanda M. Hughes, Phillip R. Acuff, and OAR Corp.					
13a. TYPE OF REPORT Final		13b. TIME COVERED FROM <u>6/89</u> TO <u>1/92</u>	14. DATE OF REPORT (Year, Month, Day) 1992, May		15. PAGE COUNT 225
16. SUPPLEMENTARY NOTATION					
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number) RTEMS, real-time, executive, heterogeneous, homogeneous, multiprocessing, 80386, microprocessor, C language, runtime, (Continued on page ii)		
FIELD	GROUP	SUB-GROUP			
19. ABSTRACT (Continue on reverse if necessary and identify by block number) This document is a detailed design manual for a real-time multiprocessor executive which provides a high performance environment for embedded military applications including such features as multitasking capabilities; homogeneous and heterogeneous multiprocessor systems; time event-driven, priority-based, preemptive scheduling; intertask communication and synchronization; responsive interrupt management; dynamic memory allocation; and a high level of user configurability. This executive, known as RTEMS (Real-Time Executive for Missile Systems), was originally developed in an effort to eliminate many of the major drawbacks of the Ada programming language. RTEMS is based on the RTEID (now ORKID) proposed standard. The code is Government owned, so no licensing fees are necessary. The executive is written using the 'C' programming language with a small amount of assembly language code. The code was developed as a linkable and/or ROMable library with the Ada programming language. Initially RTEMS was developed for the Motorola 68000 family of processors. (Continued on page ii)					
20. DISTRIBUTION / AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED		
22a. NAME OF RESPONSIBLE INDIVIDUAL Wanda M. Hughes			22b. TELEPHONE (Include Area Code) (205) 876-4484		22c. OFFICE SYMBOL AMSMI-RD-GC-S

BLOCK 18 (Cont'd): scheduling, directive, multitasking, event-driven, priority-based, preemptive, scheduling, intertask communication, synchronization, dynamic memory allocation, user configurable, kernel, embedded, semaphore, events, interrupt, regions, segments, I/O, messages, user extendable, object oriented

BLOCK 19 (Cont'd): It has been ported to the Intel 80386 and 80960 families. Other processor ports are planned for the future. This manual describes the implementation of RTEMS for the i80386 microprocessor for applications using the 'C' programming language and describes the user interface and run-time behavior. Related documents include: Real Time Executive for Missile Systems i80386 Ada Interface, Real Time Executive for Missile Systems i80386 Timing Document, and Real Time Executive for Missile Systems i80386 Assembly Interface. RTEMS documentation and code is available for the Motorola 68000 family, and the Intel 80386 and 80960 family of processors.

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	



Preface

In recent years, the cost required to develop a software product has increased significantly while the target hardware costs have decreased. Now a larger portion of money is expended in developing, using, and maintaining software. The trend in computing costs is the complete dominance of software over hardware costs. Because of this, it is necessary that formal disciplines be established to increase the probability that software is characterized by a high degree of correctness, maintainability, and portability. In addition, these disciplines must promote practices that aid in the consistent and orderly development of a software system within schedule and budgetary constraints. To be effective, these disciplines must adopt standards which channel individual software efforts toward a common goal.

The push for standards in the software development field has been met with various degrees of success. The Microprocessor Operating Systems Interfaces (MOSI) effort has experienced only limited success. As popular as the UNIX operating system has grown, the attempt to develop a standard interface definition to allow portable application development has only recently begun to produce the results needed in this area. Unfortunately, very little effort has been expended to provide standards addressing the needs of the real-time community. Several organizations have addressed this need during the past several years.

The Real Time Executive Interface Definition (RTEID) was developed by Motorola with technical input from Software Components Group. RTEID was adopted by the VMEbus International Trade Association (VITA) as a baseline

draft for their proposed standard multiprocessor, real-time executive interface, Open Real-Time Kernel Interface Definition (ORKID). These two groups are currently working together with the IEEE P1003.4 committee to insure that the functionality of their proposed standards is adopted as the real-time extensions to POSIX.

This emerging standard defines an interface for the development of real-time software to ease the writing of real-time application programs that are directly portable across multiple real-time executive implementations. This interface includes both the source code interfaces and run-time behavior as seen by a real-time application. It does not include the details of how a kernel implements these functions. The standard's goal is to serve as a complete definition of external interfaces so that application code that conforms to these interfaces will execute properly in all real-time executive environments. With the use of a standards compliant executive, routines that acquire memory blocks, create and manage message queues, establish and use semaphores, and send and receive signals need not be redeveloped for a different real-time environment as long as the new environment is compliant with the standard. Software developers need only concentrate on the hardware dependencies of the real-time system. Furthermore, most hardware dependencies for real-time applications can be localized to the device drivers.

A compliant executive provides simple and flexible real-time multiprocessing. It easily lends itself to both tightly-coupled and loosely-coupled configurations (depending on the system hardware configuration). Objects such as tasks, queues, events, signals, semaphores, and memory blocks can be designated as global objects and accessed by any task regardless of which processor the object and the accessing task reside.

The acceptance of a standard for real-time executives will produce the same advantages enjoyed from the push for UNIX standardization by AT&T's System V Interface Definition and IEEE's POSIX efforts. A compliant multiprocessing executive will allow close coupling between UNIX systems and real-time executives to provide the many benefits of the UNIX development environment to be applied to real-time software development. Together they provide the necessary laboratory environment to implement real-time, distributed, embedded systems using a wide variety of computer architectures.

A study was completed in 1988, within the Research, Development, and Engineering Center, U.S. Army Missile Command, which compared the various aspects of the Ada programming language as they related to the application of Ada code in distributed and/or multiple processing systems. Several critical conclusions were derived from the study. These conclusions have a major impact on the way the Army develops application software for embedded

applications. These impacts apply to both in-house software development and contractor developed software.

A conclusion of the analysis, which has been previously recognized by other agencies attempting to utilize Ada in a distributed or multiprocessing environment, is that the Ada programming language does not adequately support multiprocessing. Ada does provide a mechanism for multi-tasking, however this capability exists only for a single processor system. The language also does not have inherent capabilities to access global named variables, flags or program code. These critical features are essential in order for data to be shared between processors. However, these drawbacks do have *workarounds* which are sometimes awkward and defeat the intent of software maintainability goals.

Another conclusion drawn from the analysis, was that the run time executives being delivered with the Ada compilers were too slow and inefficient to be used in modern missile systems. A run time executive is the core part of the run time system code, or operating system code, that controls task scheduling, input/output management and memory management. Traditionally, whenever efficient executive (also known as kernel) code was required by the application, the user developed in-house software. This software was usually written in assembly language for optimization.

Because of this shortcoming in the Ada programming language, software developers in research and development and contractors for project managed systems, are mandated by technology to purchase and utilize off-the-shelf third party kernel code. The contractor, and eventually the Government, must pay a licensing fee for every copy of the kernel code used in an embedded system.

The main drawback to this development environment is that the Government does not own, nor has the right to modify code contained within the kernel. V&V techniques in this situation are more difficult than if the complete source code were available. Responsibility for system failures due to faulty software is yet another area to be resolved under this environment.

The Guidance and Control Directorate began a software development effort to address these problems. A project to develop an experimental run time kernel was begun that will eliminate the major drawbacks of the Ada programming language mentioned above. The Real Time Executive for Missile Systems (RTEMS) provides full capabilities for management of tasks, interrupts, time, and multiple processors in addition to those features typical of generic operating systems. The code is Government owned, so no licensing fees are necessary. The code was developed as a linkable and/or ROMable library with the Ada programming language. Initially the library code was developed on the Motorola 68000 family of processors using the 'C' programming language as the

development language. It has since been ported to the Intel 80386 family. Other language interfaces and processor families, including RISC, are planned in the future.

The final RTEMS product will be capable of handling either homogeneous or heterogeneous systems. The kernel will automatically compensate for architectural differences (byte swapping, etc.) between processors. This will allow a much easier transition from one processor family to another without a major system redesign.

Since the proposed standards are still in draft form, RTEMS cannot and does not claim compliance. However, the status of the standard is being carefully monitored to guarantee that RTEMS provides the functionality specified in the standard. Once approved, RTEMS will be made compliant.

This document is a detailed design guide for a functionally compliant real-time multiprocessor executive. It describes the user interface and run-time behavior of RTEMS.

Table of Contents

1	Overview	1
1.1	Introduction	1
1.2	Real-time Application Systems	1
1.3	Real-time Executive	2
1.4	RTEMS Application Architecture	3
1.5	RTEMS Internal Architecture	4
1.6	User Customization and Extensibility	5
1.7	Portability	5
1.8	Memory Requirements	5
1.9	Audience	6
1.10	Conventions	6
1.11	Manual Organization	7
2	Key Concepts	9
2.1	Introduction	9
2.2	Objects	9
2.3	Communication and Synchronization	11
2.4	Time	11
2.5	Memory Management	12
3	Initialization Manager	13
3.1	Introduction	13
3.2	Background	13
3.2.1	Initialization Tasks	13
3.2.2	The System Initialization Task	14
3.2.3	The Idle Task	14
3.2.4	Initialization Manager Failure	14
3.3	Operations	15
3.3.1	Initializing RTEMS	15
3.4	Directives	15
3.4.1	INIT_EXEC - Initialize RTEMS	16
4	Task Manager	17
4.1	Introduction	17
4.2	Background	18

4.2.1	Task Definition	18
4.2.2	Task Control Block	18
4.2.3	Task States	19
4.2.4	Task Priority	19
4.2.5	Task Mode	20
4.2.6	Accessing Task Arguments	21
4.2.7	Floating Point Considerations	21
4.2.8	Building an Attribute Set, Mode, or Mask	22
4.3	Operations	22
4.3.1	Creating Tasks	22
4.3.2	Obtaining Task IDs	23
4.3.3	Starting and Restarting Tasks	23
4.3.4	Suspending and Resuming Tasks	24
4.3.5	Changing Task Priority	24
4.3.6	Changing Task Mode	24
4.3.7	Notepad Locations	24
4.3.8	Task Deletion	25
4.4	Directives	25
4.4.1	T_CREATE - Create a task	26
4.4.2	T_IDENT - Get ID of a task	28
4.4.3	T_START - Start a task	29
4.4.4	T_RESTART - Restart a task	30
4.4.5	T_DELETE - Delete a task	31
4.4.6	T_SUSPEND - Suspend a task	32
4.4.7	T_RESUME - Resume a task	33
4.4.8	T_SETPRI - Set task priority	34
4.4.9	T_MODE - Change current task's mode	35
4.4.10	T_GETNOTE - Get task notepad entry	37
4.4.11	T_SETNOTE - Set task notepad entry	38
5	Interrupt Manager	39
5.1	Introduction	39
5.2	Background	39
5.2.1	i80386 Interrupt Processing	39
5.2.2	Mapping Processor Interrupt Levels	40
5.2.3	Disabling of Interrupts by RTEMS	40
5.2.4	Format of an ISR	41
5.2.5	Directives Allowed from an ISR	42
5.3	Operations	42

5.3.1	Entering an ISR	42
5.3.2	Exiting an ISR	43
5.4	Directives	43
5.4.1	I_ENTER - Enter an ISR	44
5.4.2	I_RETURN - Return from an ISR	45
6	Time Manager	47
6.1	Introduction	47
6.2	Background	48
6.2.1	Required Support	48
6.2.2	Time and Date Data Structure	48
6.2.3	Timer Types	48
6.2.4	Timeslicing	49
6.3	Operations	49
6.3.1	Announcing a Tick	49
6.3.2	Setting and Obtaining the Time	50
6.3.3	Using a Sleep Timer	50
6.3.4	Using an Event Timer	50
6.3.5	Canceling a Timer	51
6.4	Directives	51
6.4.1	TM_SET - Set system date and time	52
6.4.2	TM_GET - Get system date and time	53
6.4.3	TM_WKAFTER - Wake up after interval	54
6.4.4	TM_WKWHEN - Wake up when specified	55
6.4.5	TM_EVAFTER - Send event set after interval	56
6.4.6	TM_EVWHEN - Send event set when specified	57
6.4.7	TM_EVEVERY - Send periodic event set	58
6.4.8	TM_CANCEL - Cancel timer event	59
6.4.9	TM_TICK - Announce a clock tick	60
7	Semaphore Manager	61
7.1	Introduction	61
7.2	Background	61
7.2.1	Building an Attribute Set	62
7.3	Operations	63
7.3.1	Creating a Semaphore	63
7.3.2	Obtaining Semaphore IDs	63
7.3.3	Acquiring a Semaphore	63
7.3.4	Releasing a Semaphore	64

7.3.5	Semaphore Deletion	64
7.4	Directives	64
7.4.1	SM_CREATE - Create a semaphore	65
7.4.2	SM_IDENT - Get ID of a semaphore	67
7.4.3	SM_DELETE - Delete a semaphore	68
7.4.4	SM_P - Acquire a semaphore	69
7.4.5	SM_V - Release a semaphore	71
8	Message Manager	73
8.1	Introduction	73
8.2	Background	74
8.2.1	Messages	74
8.2.2	Message Queues	74
8.2.3	Building an Attribute Set	74
8.3	Operations	74
8.3.1	Creating a Message Queue	74
8.3.2	Obtaining Message Queue IDs	75
8.3.3	Receiving a Message	75
8.3.4	Sending a Message	75
8.3.5	Broadcasting a Message	76
8.3.6	Message Queue Deletion	76
8.4	Directives	76
8.4.1	Q_CREATE - Create a queue	77
8.4.2	Q_IDENT - Get ID of a queue	79
8.4.3	Q_DELETE - Delete a queue	80
8.4.4	Q_SEND - Put message at rear of a queue	81
8.4.5	Q_URGENT - Put message at front of a queue	82
8.4.6	Q_BROADCAST - Broadcast N messages to a queue	83
8.4.7	Q_RECEIVE - Receive message from a queue	84
8.4.8	Q_FLUSH - Flush all messages on a queue	86
9	Event Manager	87
9.1	Introduction	87
9.2	Background	87
9.2.1	Event Sets	87
9.2.2	Building an Event Set or Condition	88
9.2.3	Building a Flag	88
9.3	Operations	89
9.3.1	Sending an Event Set	89


9.3.2	Receiving an Event Set	89
9.3.3	Determining the Pending Event Set	90
9.4	Directives	90
9.4.1	EV_SEND - Send event set to a task	91
9.4.2	EV_RECEIVE - Receive event condition	92
10	Signal Manager	95
10.1	Introduction	95
10.2	Background	95
10.2.1	Definitions	95
10.2.2	A Comparison of ASRs and ISRs	96
10.2.3	Building a Signal Set	96
10.2.4	Building a Mode	97
10.3	Operations	97
10.3.1	Establishing an ASR	97
10.3.2	Sending a Signal Set	98
10.3.3	Entering an ASR	98
10.3.4	Returning from an ASR	98
10.3.5	Format of an ASR	98
10.4	Directives	99
10.4.1	AS_CATCH - Establish an ASR	100
10.4.2	AS_SEND - Send signal set to a task	101
10.4.3	AS_ENTER - Enter an ASR	102
10.4.4	AS_RETURN - Return from an ASR	103
11	Partition Manager	105
11.1	Introduction	105
11.2	Background	105
11.2.1	Definitions	105
11.2.2	Building an Attribute Set	106
11.3	Operations	106
11.3.1	Creating a Partition	106
11.3.2	Obtaining Partition IDs	106
11.3.3	Acquiring a Buffer	106
11.3.4	Releasing a Buffer	107
11.3.5	Deleting a Partition	107
11.4	Directives	107
11.4.1	PT_CREATE - Create a partition	108
11.4.2	PT_IDENT - Get ID of a partition	110

11.4.3	PT_DELETE - Delete a partition	111
11.4.4	PT_GETBUF - Get buffer from a partition	112
11.4.5	PT_RETBUF - Return buffer to a partition	113
12	Region Manager	115
12.1	Introduction	115
12.2	Background	115
12.2.1	Definitions	115
12.2.2	Building an Attribute Set	116
12.3	Operations	116
12.3.1	Creating a Region	116
12.3.2	Obtaining Region IDs	117
12.3.3	Acquiring a Segment	117
12.3.4	Releasing a Segment	117
12.3.5	Deleting a Region	118
12.4	Directives	118
12.4.1	RN_CREATE - Create a region	119
12.4.2	RN_IDENT - Get ID of a region	121
12.4.3	RN_DELETE - Delete a region	122
12.4.4	RN_GETSEG - Get segment from a region	123
12.4.5	RN_RETSEG - Return segment to a region	125
13	Dual-Ported Memory Manager	127
13.1	Introduction	127
13.2	Background	127
13.3	Operations	128
13.3.1	Creating a Port	128
13.3.2	Obtaining Port IDs	128
13.3.3	Converting an Address	128
13.3.4	Deleting a DPMA Port	129
13.4	Directives	129
13.4.1	DP_CREATE - Create a port	130
13.4.2	DP_IDENT - Get ID of a port	131
13.4.3	DP_DELETE - Delete a port	132
13.4.4	DP_2INTERNAL - Convert external to internal address	133
13.4.5	DP_2EXTERNAL - Convert internal to external address	134

14	I/O Manager	135
14.1	Introduction	135
14.2	Background	135
14.2.1	Device Driver Table	135
14.2.2	Major and Minor Device Numbers	136
14.2.3	Device Driver Environment	136
14.2.4	Device Driver Interface	137
14.2.5	Device Driver Initialization	137
14.3	Operations	138
14.4	Directives	138
14.4.1	DE_INIT - Initialize a device driver	139
14.4.2	DE_OPEN - Open a device	140
14.4.3	DE_CLOSE - Close a device	141
14.4.4	DE_READ - Read from a device	142
14.4.5	DE_WRITE - Write to a device	143
14.4.6	DE_CNTRL - Special device services	144
15	Fatal Error Manager	145
15.1	Introduction	145
15.2	Background	145
15.3	Operations	146
15.4	Announcing a Fatal Error	146
15.5	Directives	146
15.5.1	K_FATAL - Invoke the fatal error handler	147
16	Scheduling Concepts	149
16.1	Introduction	149
16.2	Scheduling Mechanisms	150
16.2.1	Task Priority	150
16.2.2	Preemption	150
16.2.3	Timeslicing	151
16.2.4	Manual Round-Robin	151
16.2.5	Dispatching Tasks	151
16.3	Task State Transitions	152
17	Board Support Packages	157
17.1	Introduction	157
17.2	System Reset and Initialization	157

17.2.1	SEGMENT USAGE	159
17.2.2	Allowing for Interrupt Stack Usage	160
18	User Extensions	161
18.1	Introduction	161
18.2	TCREATE Extension	162
18.3	TSTART Extension	162
18.4	TRESTART Extension	162
18.5	TDELETE Extension	163
18.6	TSWITCH Extension	163
18.7	TASKEXITTED Error Extension	163
18.8	FATAL Error Extension	164
18.9	TCB Extension	164
19	Configuring a System	165
19.1	Configuration Table	165
19.2	Initialization Task Table	167
19.3	Driver Address Table	168
19.4	User Extensions Table	169
19.5	Multiprocessor Configuration Table	171
19.6	Multiprocessor Communications Interface Table	172
19.7	Determining Memory Requirements	173
19.7.1	Sizing the RTEMS RAM Workspace	174
20	Multiprocessing Manager	175
20.1	Introduction	175
20.2	Background	176
20.2.1	Nodes	176
20.2.2	Global Objects	176
20.2.3	Global Object Table	177
20.2.4	Remote Operations	177
20.2.5	Proxies	178
20.2.6	Multiprocessor Configuration Table	179
20.3	Multiprocessor Communications Interface Layer	179
20.3.1	INIT	180
20.3.2	GETPKT	180
20.3.3	RETPKT	181
20.3.4	RECEIVE	181

20.3.5	SEND	181
20.3.6	Supporting Heterogeneous Environments	182
20.4	Operations	183
20.4.1	Announcing a Packet	183
20.5	Directives	183
20.5.1	MP_ANNOUNCE - Announce the arrival of a packet	184
A	Memory Requirements	1
A.1	Data Space Requirements	1
A.2	Minimum and Maximum Code Space Requirements	1
A.3	Code Space Worksheet	2
A.4	RTEMS RAM Workspace Worksheet	3
B	DIRECTIVE STATUS CODES	1
C	HEADER FILES	1
C.1	Header File Usage	1
C.2	rtems.h	1
D	EXAMPLE APPLICATION	1
E	GLOSSARY	1



Overview

1.1 Introduction

RTEMS, *Real-Time Executive for Missile Systems*, is a real-time executive (kernel) which provides a high performance environment for embedded military applications including the following features:

- *multitasking capabilities*
- *homogeneous and heterogeneous multiprocessor systems*
- *event-driven, priority-based, preemptive scheduling*
- *intertask communication and synchronization*
- *responsive interrupt management*
- *dynamic memory allocation*
- *high level of user configurability*

This manual describes the implementation of **RTEMS** for the target microprocessor for applications using the **C** programming language.

1.2 Real-time Application Systems

Real-time application systems are a special class of computer applications. They have a complex set of characteristics that distinguish them from other software problems. Generally, they must adhere to more rigorous requirements. The correctness of the system depends not only on the results of computations, but also on the time at which the results are produced. The most important and complex characteristic of real-time application systems is that they must receive and respond to a set of external stimuli within rigid and critical time constraints.

Systems can be buried by an avalanche of interdependent, asynchronous or cyclical event streams.

Another distinguishing requirement of real-time application systems is the ability to coordinate or manage a large number of concurrent activities. Since software is a synchronous entity, this presents special problems. One instruction follows another in a repeating synchronous cycle. Even though mechanisms have been developed to allow for the processing of external asynchronous events, the software design efforts required to process and manage these events and tasks are growing more complicated.

The design process is complicated further by spreading this activity over a set of processors instead of a single processor. The challenges associated with designing and building real-time application systems become very complex when multiple processors are involved. New requirements such as interprocessor communication channels and global resources that must be shared between competing processors are introduced. The ramifications of multiple processors complicate each and every characteristic of a real-time system.

1.3 Real-time Executive

Fortunately, real-time operating systems or real-time executives serve as a cornerstone on which to build the application system. A real-time multitasking executive allows an application to be cast into a set of logical, autonomous processes or tasks which become quite manageable. Each task is internally synchronous, but different tasks execute independently, resulting in an asynchronous processing stream. Tasks can be dynamically paused for many reasons resulting in a different task being allowed to execute for a period of time. The executive also provides an interface to other system components such as interrupt handlers and device drivers. System components may request the executive to allocate and coordinate resources, and to wait for and trigger synchronizing conditions. The executive system calls effectively extend the CPU instruction set to support efficient multitasking. By causing tasks to travel through well-defined state transitions, system calls permit an application to demand-switch between tasks in response to real-time events.

By proper grouping of responses to stimuli into separate tasks, a system can now asynchronously switch between independent streams of execution, directly responding to external stimuli as they occur. This allows the system design to meet critical performance specifications which are typically measured by guaranteed response time and transaction throughput. The multiprocessor extensions of RTEMS provide the features necessary to manage the extra requirements introduced by a system distributed across several processors. It removes the physical barriers of processor boundaries from the world of the

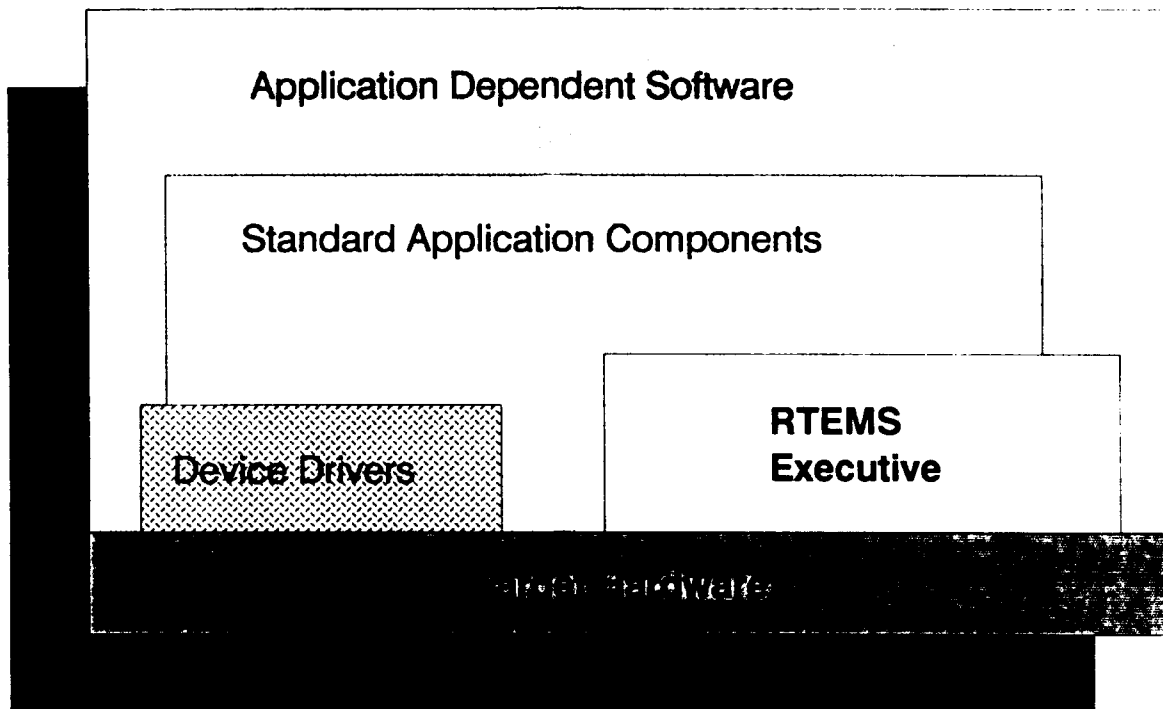


Figure 1-1 RTEMS Application Architecture

system designer, enabling more critical aspects of the system to receive the required attention. Such a system, based on an efficient real-time, multiprocessor executive, is a more realistic model of the outside world or environment for which it is designed. As a result, the system will always be more logical, efficient, and reliable.

By using the directives provided by RTEMS, the real-time applications developer is freed from the problem of controlling and synchronizing multiple tasks and processors. In addition, one need not develop, test, debug, and document routines to manage memory, pass messages, or provide mutual exclusion. The developer is then able to concentrate solely on the application. By using standard software components, the time and cost required to develop sophisticated real-time applications is significantly reduced.

1.4 RTEMS Application Architecture

One important design goal of RTEMS was to provide a bridge between two critical layers of typical real-time systems. As shown in Figure 1-1, RTEMS serves as a buffer between the project dependent application code and the target hardware. Most hardware dependencies for real-time applications can be localized to the low level device drivers. The RTEMS I/O interface manager provides an efficient tool for incorporating these hardware dependencies into

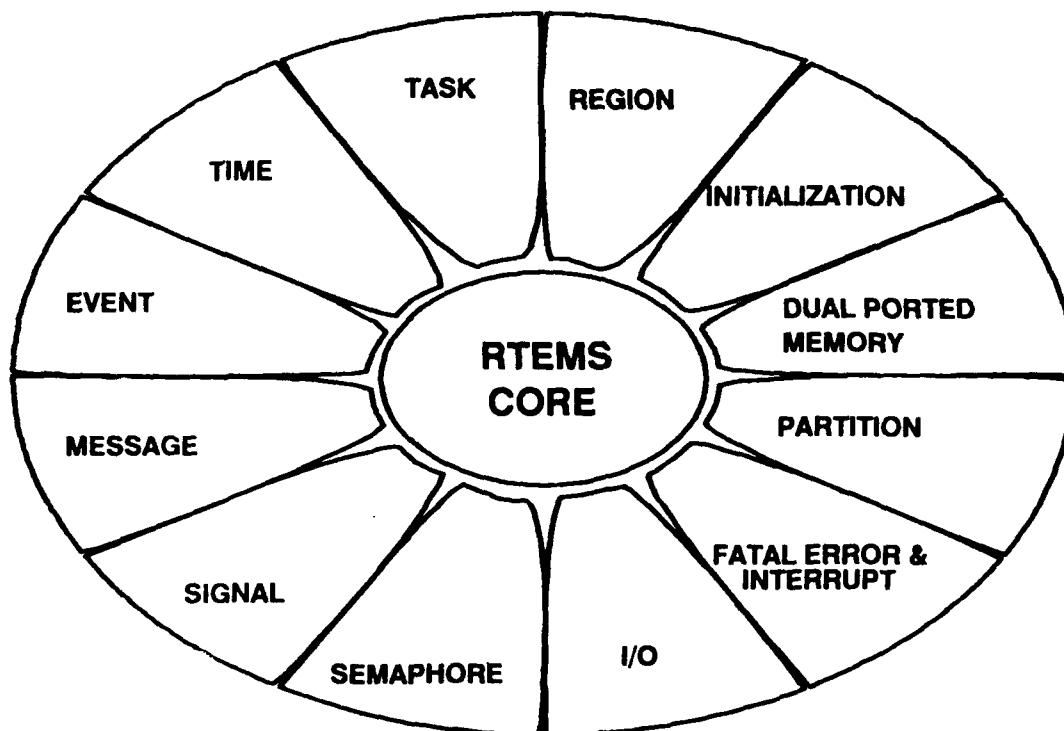


Figure 1-2 RTEMS Internal Architecture

the system while simultaneously providing a general mechanism to the application code that accesses them. A well designed real-time system can benefit from this architecture by building a rich library of standard application components which can be used repeatedly in other real-time projects.

1.5 RTEMS Internal Architecture

As illustrated in Figure 1-2, RTEMS can be viewed as a set of components that work in harmony to provide a set of services to a real-time application system. The executive interface presented to the application is formed by grouping directives into logical sets called **resource managers**. Functions utilized by multiple managers such as scheduling, dispatching, and object management are provided in the **executive core**. Together these components provide a powerful run time environment that promotes the development of efficient real-time application systems. Subsequent chapters present a detailed description of the capabilities provided by each of the following RTEMS managers:

- *initialization*
- *task*
- *interrupt*
- *signal*
- *partition*
- *region*

- *time*
- *semaphore*
- *message*
- *event*
- *dual ported memory*
- *I/O*
- *fatal error*

The **C Interface Library** component is a collection of routines which allow C application programs to invoke **RTEMS** directives. This extends the standard C language to include the real-time features provided by **RTEMS** without modifying the C compiler. Only three **RTEMS** directives, related to interrupts (*i_enter* and *i_return*) and asynchronous signals (*as_return*), cannot be accessed from C due to their processor dependent nature.

1.6 User Customization and Extensibility

As thirty-two bit microprocessors have decreased in cost, they have become increasingly common in a variety of embedded systems. A wide range of custom and general-purpose processor boards are based on various thirty-two bit processors. **RTEMS** was designed to make no assumptions concerning the characteristics of individual microprocessor families or of specific support hardware. In addition, **RTEMS** allows the system developer a high degree of freedom in customizing and extending its features.

RTEMS assumes the existence of a supported microprocessor and sufficient memory for both **RTEMS** and the real-time application. Board dependent components such as clocks, interrupt controllers, or I/O devices can be easily integrated with **RTEMS**. The customization and extensibility features allow **RTEMS** to efficiently support as many environments as possible.

1.7 Portability

The issue of portability was the major factor in the creation of **RTEMS**. Since **RTEMS** is designed to isolate the hardware dependencies in the specific board support packages, the real-time application should be easily ported to any other processor. The use of **RTEMS** in conjunction with the **C Interface Library** allows the development of real-time applications which can be completely independent of a particular microprocessor architecture.

1.8 Memory Requirements

Since memory is a critical resource in many real-time embedded systems, **RTEMS** was specifically designed to allow unused managers to be excluded from the run-time environment. This allows the application designer the flexibility

to tailor **RTEMS** to most efficiently meet system requirements while still satisfying even the most stringent memory constraints. As result, the size of the **RTEMS** executive is application dependent. Appendix A provides a worksheet for calculating the memory requirements of a custom **RTEMS** run-time environment. The following managers may be optionally excluded:

- *signal*
- *region*
- *dual ported memory*
- *I/O*
- *event*
- *partition*
- *time*
- *semaphore*
- *message*

RTEMS utilizes memory for both code and data space. Although **RTEMS**' data space must be in RAM, its code space can be located in either ROM or RAM.

1.9 Audience

This manual was written for experienced real-time software developers. Although some background is provided, it is assumed that the reader is familiar with the concepts of task management as well as intertask communication and synchronization. Since directives, user related structures, and examples are presented in C, a basic understanding of the C programming language is required. A working knowledge of the target processor is helpful in understanding some of **RTEMS**' features. A thorough understanding of the executive cannot be obtained without studying the entire manual because many of **RTEMS**' concepts and features are interrelated. Experienced **RTEMS** users will find that the manual organization facilitates its use as a reference document.

1.10 Conventions

The following conventions are used in this manual:

- *Significant words or phrases as well as all directive names are printed in bold type.*
- *Items in bold capital letters are constants defined by **RTEMS**. Each language interface provided by **RTEMS** includes a file containing the standard set of constants, data types, and structure/record definitions which can be incorporated into the user application.*

- *A number of type definitions are provided by RTEMS and can be found in **rtems.h**.*
- *The characters "0x" preceding a number indicates that the number is in hexadecimal format. Any other numbers are assumed to be in decimal format.*
- *The ampersand character (&) preceding a symbol indicates that the address of the symbol is passed to the called routine instead of the value itself.*

1.11 Manual Organization

This first chapter has presented the introductory and background material for the RTEMS executive. The remaining chapters of this manual present a detailed description of RTEMS and the environment, including run time behavior, it creates for the user.

Chapter 2: Key Concepts: presents an introduction to the ideas which are common across multiple RTEMS managers.

Chapters 3 - 15: Managers and Directives: these chapters provide a detailed discussion of each RTEMS manager and the directives which it provides. The presentation format for each directive includes the following sections:


- *Calling sequence*
- *Input parameters*
- *Output parameters*
- *Directive status codes*
- *General description*
- *Notes*

Chapter 16: Scheduling Concepts: details the RTEMS scheduling algorithm and task state transitions.

Chapter 17: Board Support Packages: defines the functionality required of user-supplied board support packages.

Chapter 18: User Extensions: shows the user how to extend RTEMS to incorporate custom features .

- Chapter 19:** **Configuring a System:** details the process by which one tailors **RTEMS** for a particular single-processor or multiprocessor application.
- Chapter 20:** **Multiprocessing:** presents a conceptual overview of the multiprocessing capabilities provided by **RTEMS** and describes the **Multiprocessing Communications Interface Layer**.
- Appendix A:** **RTEMS Memory Requirements Worksheets:** provides a worksheet for calculating the code and data space requirements for a custom configuration of **RTEMS**.
- Appendix B:** **Directive Status Codes:** provides a definition of each of the directive status codes referenced in this manual. These definitions are also provided in the user include file **rtems.h**.
- Appendix C:** **Definition File:** describes and lists the **RTEMS** provided constants, data types, and structure/record definitions.
- Appendix D:** **Example RTEMS Application:** provides a template for simple **RTEMS** applications.
- Appendix E:** **Glossary:** defines terms used throughout this manual.



Key Concepts

2.1 Introduction

The facilities provided by **RTEMS** are built upon a foundation of very powerful concepts. These concepts must be understood before the application developer can efficiently utilize **RTEMS**. The purpose of this chapter is to familiarize one with these concepts.

2.2 Objects

RTEMS provides directives which can be used to dynamically create, delete, and manipulate a set of predefined object types. These types include tasks, message queues, semaphores, memory regions, memory partitions, and timers. The object-oriented nature of **RTEMS** encourages the creation of modular applications built upon re-usable "building block" routines.

All objects are created on the local node as required by the application and have an **RTEMS** assigned ID. All objects except timers have a user-assigned name. Although a relationship exists between an object's name and its **RTEMS** assigned ID, the name and ID are not identical. Object names are completely arbitrary and selected by the user as a meaningful "tag" which may commonly reflect the object's use in the application. Conversely, object IDs are designed to facilitate efficient object manipulation by the executive.

An object **name** is an unsigned thirty-two bit entity associated with the object by the user. Although not required by **RTEMS**, object names are typically composed of four ASCII characters which help identify that object. For example, a task which causes a light to blink might be called "LITE". On the other hand, if an application requires one-hundred tasks, it would be difficult to assign meaningful ASCII names to each task. A more convenient approach would be to name them the binary values one through one-hundred, respectively.




Figure 2-1 Object ID Composition

An **object ID** is a unique unsigned thirty-two bit entity composed of two parts. The most significant sixteen bits are the number of the node on which this object was created. The node number is always one (1) in a single processor system. The least significant sixteen bits form an identifier within a particular object type. This identifier, called the **object index**, ranges in value from 1 to the maximum number of objects configured for this object type.

The two components of an object ID make it possible to quickly locate any object in even the most complicated multiprocessor system. Object ID's are associated with an object by **RTEMS** when the object is created and the corresponding ID is returned by the appropriate object create directive. The object ID is required as input to all directives involving objects, except those which create an object or obtain the ID of an object.

The object identification directives can be used to dynamically obtain a particular object's ID given its name. This mapping is accomplished by searching the name table associated with this object type. If the name is non-unique, then the ID associated with the first occurrence of the name will be returned to the application. Since object IDs are returned when the object is created, the object identification directives are not necessary in a properly designed single processor application.

An **object control block** is a data structure defined by **RTEMS** which contains the information necessary to manage a particular object type. For efficiency reasons, the format of each object type's control block is different. However, many of the fields are similar in function. The number of each type of control block is application dependent and determined by the values specified in the user's **Configuration Table**. An object control block is allocated at object create time and freed when the object is deleted. With the exception of user extension routines, object control blocks are not directly manipulated by user applications.

2.3 Communication and Synchronization

In real-time multitasking applications, the ability for cooperating execution threads to communicate and synchronize with each other is imperative. A real-time executive should provide an application with the following capabilities:

- *Data transfer between cooperating tasks*
- *Data transfer between tasks and ISRs*
- *Synchronization of cooperating tasks*
- *Synchronization of tasks and ISRs*

Most RTEMS managers can be used to provide some form of communication and/or synchronization. However, managers dedicated specifically to communication and synchronization provide well established mechanisms which directly map to the application's varying needs. This level of flexibility allows the application designer to match the features of a particular manager with the complexity of communication and synchronization required. The following managers were specifically designed for communication and synchronization:

- *Semaphore*
- *Event*
- *Message*
- *Signal*

The semaphore manager supports mutual exclusion involving the synchronization of access to one or more shared user resources. The message manager supports both communication and synchronization, while the event manager primarily provides a high performance synchronization mechanism. The signal manager supports only asynchronous communication and is typically used for exception handling.

2.4 Time

The development of responsive real-time applications requires an understanding of how RTEMS maintains and supports time-related operations. The basic unit of time in RTEMS is known as a **tick**. The frequency of clock ticks is completely application dependent and determines the granularity and accuracy of all interval and calendar time operations.

By tracking time in units of ticks, RTEMS is capable of supporting interval timing functions such as task delays, timeouts, timeslicing, and the delayed posting of events. An **interval** is defined as a number of ticks relative to the

current time. For example, when a task delays for an interval of ten ticks, it is implied that the task will not execute until ten clock ticks have occurred.

A characteristic of interval timing is that the actual interval period may be a fraction of a tick less than the interval requested. This occurs because the time at which the delay timer is set up occurs at some time between two clock ticks. Therefore, the first countdown tick occurs in less than the complete time interval for a tick. This can be a problem if the clock granularity is large.

Interval timing is not sufficient for the many applications which require that time be kept in wall time or true calendar form. Consequently, RTEMS maintains the current date and time. This allows selected time operations to be scheduled at an actual calendar date and time. For example, a task could request to delay until midnight on New Year's Eve before lowering the ball at Times Square.


Obviously, the time manager's directives cannot operate without some external mechanism which provides a periodic clock tick. This clock tick is typically provided by a real time clock or counter/timer device.

2.5 Memory Management

RTEMS memory management facilities can be grouped into two classes: dynamic memory allocation and address translation. Dynamic memory allocation is required by applications whose memory requirements vary through the application's course of execution. Address translation is needed by applications which share memory with another CPU or an intelligent Input/Output processor. The following RTEMS managers provide facilities to manage memory:

- *Region*
- *Dual Ported Memory*
- *Partition*

RTEMS memory management features allow an application to create simple memory pools of fixed size buffers and/or more complex memory pools of variable size segments. The partition manager provides directives to manage and maintain pools of fixed size entities such as resource control blocks. Alternatively, the region manager provides a more general purpose memory allocation scheme that supports variable size blocks of memory which are dynamically obtained and freed by the application. The dual-ported memory manager provides executive support for address translation between internal and external dual-ported RAM address space.



Initialization Manager

3.1 Introduction

The **initialization manager** is responsible for initiating **RTEMS**, creating and starting all configured initialization tasks, and for invoking the initialization routine for each user-supplied device driver. In a multiprocessor configuration, this manager also initializes the interprocessor communications layer. The directive provided by the initialization manager is:

Name	Directive Description
<code>init_exec</code>	Initialize RTEMS

3.2 Background

3.2.1 Initialization Tasks

Initialization task(s) are the mechanism by which **RTEMS** transfers initial control to the user's application. Initialization tasks differ from other application tasks in that they are defined in the **Initialization Task Table** and automatically created and started by **RTEMS** as part of its initialization sequence. Since the initialization tasks are scheduled using the same algorithm as all other **RTEMS** tasks, they must be configured at a priority and mode which will insure that they will complete execution before other application tasks execute. Although there is no upper limit on the number of initialization tasks, an application is required to define at least one.

A typical initialization task will create and start the static set of application tasks. It may also create any other objects used by the application. Initialization tasks which only perform initialization should delete themselves upon completion to free resources for other tasks. Initialization tasks may transform themselves into a "normal" application task. This transformation typically involves changing priority and execution mode. **RTEMS does not automatically delete the initialization tasks.**

3.2.2 The System Initialization Task

The **System Initialization Task** is responsible for initializing all device drivers. As a result, this task has a higher priority than all other application tasks to insure that no application tasks executes until all device drivers are initialized. After device initialization in a single processor system, this task will delete itself.

In multiprocessor configurations, the **System Initialization Task** does not delete itself after initializing the device drivers. Instead it transforms itself into the **Multiprocessing Server** which initializes the **Multiprocessor Communications Interface** layer, verifies multiprocessor system consistency, and processes all requests from remote nodes.

3.2.3 The Idle Task

The **Idle Task** is the lowest priority task in all systems and executes only when no other task is ready to execute. This task consists of an infinite loop and will be preempted when any other task is made ready to execute.

3.2.4 Initialization Manager Failure

The **k_fatal** directive will be called from **init_exec** for any of the following reasons:

- *If no user initialization tasks are configured. At least one initialization task must be configured to allow RTEMS to pass control to the application at the end of the executive initialization sequence.*
- *If the starting address of the RTEMS RAM Workspace, supplied by the application in the Configuration Table, is not aligned on a four-byte boundary.*
- *If the size of the RTEMS RAM Workspace is not large enough to initialize and configure the system.*

- *If multiprocessing is configured and the node entry in the Multiprocessor Configuration Table is not between one and the max_nodes entry.*
- *If any of the user initialization tasks cannot be created or started successfully.*

3.3 Operations

3.3.1 Initializing RTEMS

The `init_exec` directive is called by the board support package at the completion of its initialization sequence. RTEMS assumes that the board support package successfully completed its initialization activities. The `init_exec` directive completes the initialization sequence by performing the following actions:

- *Initialize internal RTEMS variables;*
- *Allocate system resources;*
- *Create and start the System Initialization Task;*
- *Create and start the Idle Task;*
- *Create and start the user initialization task(s); and*
- *Initiate multitasking.*

This directive **MUST** be called before any other RTEMS directives. The effect of calling any RTEMS directives before `init_exec` is unpredictable. Many of RTEMS actions during initialization are based upon the contents of the Configuration Table. For more information regarding the format and contents of this table, please refer to the chapter **Configuring a Single Processor System**.

The final step in the initialization sequence is the initiation of multitasking. When the scheduler and dispatcher are enabled, the highest priority, ready task will be dispatched to run. Control will not be returned to the board support package after multitasking is enabled.

3.4 Directives

This section details the initialization manager's directives. A subsection is dedicated to each of this manager's directives and describes the calling sequence, related constants, usage, and status codes.

3.4.1 INIT_EXEC - Initialize RTEMS

CALLING SEQUENCE:

```
void init_exec( conftbl )
```

INPUT:

```
config_table *conftbl;      /* configuration table pointer */
```

OUTPUT: NONE

DIRECTIVE STATUS CODES: NONE

DESCRIPTION:

This directive is called when the board support package has completed its initialization to allow **RTEMS** to initialize the application environment based upon the information in the **Configuration Table**.

NOTES:

This directive **MUST** be the first **RTEMS** directive called and it **DOES NOT RETURN** to the caller.

This directive causes all nodes in the system to verify that certain configuration parameters are the same as those of the local node. If an inconsistency is detected, then a fatal error is generated.



Task Manager

4.1 Introduction

The **task manager** provides a comprehensive set of directives to create, delete, and administer tasks. The directives provided by the task manager are:

Name	Directive Description
t_create	Create a task
t_ident	Get ID of a task
t_start	Start a task
t_restart	Restart a task
t_delete	Delete a task
t_suspend	Suspend a task
t_resume	Resume a task
t_setpri	Set task priority
t_mode	Change current task's mode
t_getnote	Get task notepad entry
t_setnote	Set task notepad entry

4.2 Background

4.2.1 Task Definition

Many definitions of a task have been proposed in computer literature. Unfortunately, none of these definitions encompasses all facets of the concept in a manner which is operating system independent. Several of the more common definitions are provided to enable each user to select a definition which best matches their own experience and understanding of the task concept:

- *the "dispatchable" unit.*
- *the entity to which the processor is allocated.*
- *the atomic unit of a real-time, multiprocessor system.*
- *single threads of execution which concurrently compete for resources.*
- *a sequence of closely related computations which can execute concurrently with other computational sequences.*

From RTEMS' perspective, a task is the smallest thread of execution which can compete on its own for system resources. A task is manifested by the existence of a task control block (TCB).

4.2.2 Task Control Block

The Task Control Block (TCB) is an RTEMS defined data structure which contains all the information that is pertinent to the execution of a task. During system initialization, RTEMS reserves a TCB for each task configured. A TCB is allocated upon creation of the task and is returned to the TCB free list upon deletion of the task.

The TCB's elements are modified as a result of system calls made by the application in response to external and internal stimuli. TCBs are the only RTEMS internal data structure that can be accessed by an application via user extension routines. The TCB contains a task's name, ID, current priority, current and starting states, execution mode, set of notepad locations, TCB user extension pointer, scheduling control structures, as well as data required by a blocked task.

A task's context is stored in the TCB when a task switch occurs. When the task regains control of the processor, its context is restored from the TCB. When a

task is restarted, the initial state of the task is restored from the starting context area in the task's TCB.

4.2.3 Task States

A task may exist in one of the following five states:

Task State	Description
executing	Currently scheduled to the CPU
ready	May be scheduled to the CPU
blocked	Unable to be scheduled to the CPU
dormant	Created task that is not started
non-existent	Uncreated or deleted task

Figure 4-1 Task States

An active task may occupy the executing, ready, blocked or dormant state, otherwise the task is considered non-existent. One or more tasks may be active in the system simultaneously. Multiple tasks communicate, synchronize, and compete for system resources with each other via system calls. The multiple tasks appear to execute in parallel, but actually each is dispatched to the CPU for periods of time determined by the **RTEMS** scheduling algorithm. The scheduling of a task is based on its current state and priority.

4.2.4 Task Priority

A task's priority determines its importance in relation to the other tasks executing on the same processor. **RTEMS** supports 255 levels of priority ranging from 1 to 255. Tasks of numerically smaller priority values are more important tasks than tasks of numerically larger priority values. For example, a task at priority level 5 is of higher privilege than a task at priority level 10. There is no limit to the number of tasks assigned to the same priority.

Each task has a priority associated with it at all times. The initial value of this priority is assigned at task creation time. The priority of a task may be changed at any subsequent time.

Priorities are used by the scheduler to determine which ready task will be allowed to execute. In general, the higher the priority of a task, the more likely it is to receive processor execution time.

4.2.5 Task Mode

A task's mode is a combination of the following four components:

- *preemption*
- *ASR processing*
- *timeslicing*
- *interrupt level*

It is used to modify RTEMS' scheduling process and to alter the execution environment of the task.

The preemption component allows a task to determine when control of the processor is relinquished. If preemption is disabled (**NOPREEMPT**), the task will retain control of the processor as long as it is in the ready state -- even if a higher priority task is made ready. If preemption is enabled (**PREEMPT**) and a higher priority task is made ready, then the processor will be taken away from the current task immediately and given to the higher priority task.

The timeslicing component is used by the RTEMS scheduler to determine how the processor is allocated to tasks of equal priority. If timeslicing is enabled (**TSLICE**), then RTEMS will limit the amount of time the task can execute before the processor is allocated to another ready task of equal priority. The length of the timeslice is application dependent and specified in the Configuration Table. If timeslicing is disabled (**NOTSLICE**), then the task will be allowed to execute until a task of higher priority is made ready. If **NOPREEMPT** is selected, then the timeslicing component is ignored by the scheduler.

The asynchronous signal processing component is used to determine when received signals are to be processed by the task. If signal processing is enabled (**ASR**), then signals sent to the task will be processed the next time the task executes. If signal processing is disabled (**NOASR**), then all signals received by the task will remain posted until signal processing is enabled. This component affects only tasks which have established a routine to process asynchronous signals.

The interrupt level component is used to determine which interrupts will be enabled when the task is executing. **INTR(n)** specifies that the task will execute at interrupt level **n**.

CONSTANT	DESCRIPTION	DEFAULT
PREEMPT	enable preemption	*
NOPREEMPT	disable preemption	
NOTSLICE	disable timeslicing	*
TSlice	enable timeslicing	
ASR	enable ASR processing	*
NOASR	disable ASR processing	
INTR(0)	enable all interrupts	*
INTR(n)	execute at interrupt level n	

Figure 4-2 Task Mode Constants

4.2.6 Accessing Task Arguments

All **RTEMS** tasks are invoked with a single argument which is specified when they are started or restarted. The argument is commonly used to communicate some startup information to the task. The simplest manner in which to define a task which accesses its argument is:

```
task user_task( arg )
unsigned32 arg;
```

Application tasks requiring more information may view the single argument as a pointer to a parameter block. A task utilizing the argument in this manner may be defined as follows:

```
task user_task( arg_ptr )
struct user_task_args *arg_ptr;
```

where the structure, **user_task_args** is an application defined entity.

4.2.7 Floating Point Considerations

Creating a task with the **FP** flag results in additional memory being allocated for the TCB to store the state of the numeric coprocessor during task switches. This additional memory is not allocated for **NOFP** tasks. Saving and restoring the context of a **FP** task takes longer than that of a **NOFP** task because of the relatively large amount of time required for the numeric coprocessor to save or restore its computational state.

Since **RTEMS** was designed specifically for embedded missile applications which are floating point intensive, the executive is optimized to avoid unnecessarily saving and restoring the state of the numeric coprocessor. The state of the numeric coprocessor is only saved when an **FP** task is dispatched and that task was not the last task to utilize the coprocessor. In a system with only one **FP** task, the state of the numeric coprocessor will never be saved or restored.

Although the overhead imposed by **FP** tasks is minimal, some applications may wish to completely avoid the overhead associated with **FP** tasks and still utilize a numeric coprocessor. By preventing a task from being preempted while performing a sequence of floating point operations, a **NOFP** task can utilize the numeric coprocessor without incurring the overhead of a **FP** context switch. However, if this approach is taken by the application designer, **NO** tasks should be created as **FP** tasks.

4.2.8 Building an Attribute Set, Mode, or Mask

In general, an attribute set, mode, or mask is built by a bitwise OR of the desired options. The set of valid options is provided in the description of the appropriate directive. An option listed as a default is not required to appear in the option OR list, although it is a good programming practice to specify default options. If all defaults are desired, the option **DEFAULTS** should be specified on this call.

This example demonstrates the **attr** parameter needed to create a local task which utilizes the numeric coprocessor. The **attr** parameter could be **FP** or **LOCAL|FP**. The **attr** parameter can be set to **FP** because **LOCAL** is the default for all created tasks. If the task were global and used the numeric coprocessor, then the **attr** parameter would be **GLOBAL|FP**.

The following example will demonstrate the **mode** and **mask** parameters used with the **t_mode** directive to place a task at interrupt level 3 and make it non-preemptible. The **mode** should be set to **INTR(3)|NOPREEMPT** to indicate the desired preemption mode and interrupt level, while the **mask** parameter should be set to **INTRMODE|PREEMPTMODE** to indicate that the calling task's interrupt level and preemption mode are being altered.

4.3 Operations

4.3.1 Creating Tasks

The **t_create** directive creates a task by allocating a task control block, assigning the task a user-specified name, allocating it a stack and floating point context

area, setting a user-specified initial priority, and assigning it a task ID. Newly created tasks are initially placed in the dormant state. All RTEMS tasks execute in the most privileged mode of the processor.

4.3.2 Obtaining Task IDs

When a task is created, RTEMS generates a unique task ID and assigns it to the created task until it is deleted. The task ID may be obtained by either of two methods. First, as the result of an invocation of the `t_create` directive, the task ID is stored in a user provided location. Second, the task ID may be obtained later using the `t_ident` directive. The task ID is used by other directives to manipulate this task.

4.3.3 Starting and Restarting Tasks

The `t_start` directive is used to place a dormant task in the ready state. This enables the task to compete, based on its current priority, for the processor and other system resources. Any actions, such as suspension or change of priority, performed on a task prior to starting it are nullified when the task is started.

With the `t_start` directive the user specifies the task's starting address, initial execution mode, and argument. The argument is used to communicate some startup information to the task. As part of this directive RTEMS initializes the task's stack based upon the task's initial execution mode and start address. The initialized stack will contain the specified argument and the entry address of a routine which is invoked if the task exits instead of deleting itself. The initial stack of a task is shown in Figure 4-1.




Figure 4-3 Initial Task Stack

The `t_restart` directive restarts a task at its initial starting address with its original priority and execution mode, but with a different argument. The new argument may be used to distinguish between the original invocation of the task and subsequent invocations. The task's stack and control block are modified to reflect their original creation values. Although references to resources that have been requested are cleared, resources allocated by the task are NOT automatically returned to RTEMS. A task cannot be restarted unless it has

previously been started (i.e. dormant tasks cannot be restarted). All restarted tasks are placed in the ready state.

4.3.4 Suspending and Resuming Tasks

The **t_suspend** directive is used to place either the caller or another task into a suspended state. The task remains suspended until a **t_resume** directive is issued. This implies that a task may be suspended as well as blocked waiting either to acquire a resource or for the expiration of a timer.

The **t_resume** directive is used to remove another task from the suspended state. If the task is not also blocked, resuming it will place it in the ready state, allowing it to once again compete for the processor and resources. If the task was blocked as well as suspended, this directive clears the suspension and leaves the task in the blocked state.

4.3.5 Changing Task Priority

The **t_setpri** directive is used to obtain or change the current priority of either the calling task or another task. If the new priority requested is **CURRENT** or the task's actual priority, then the current priority will be returned and the task's priority will remain unchanged. If the task's priority is altered, then the task will be scheduled according to its new priority.

The **t_restart** directive resets the priority of a task to its original value.

4.3.6 Changing Task Mode

The **t_mode** directive is used to obtain or change the current execution mode of the calling task. A task's execution mode is used to enable preemption, timeslicing, ASR processing, and to set the task's interrupt level.

The **t_restart** directive resets the mode of a task to its original value.

4.3.7 Notepad Locations

RTEMS provides sixteen notepad locations for each task. Each notepad location may contain a note consisting of four bytes of information. **RTEMS** provides two directives, **t_setnote** and **t_getnote**, that enable a user to access and change the notepad locations. The **t_setnote** directive enables the user to set a task's notepad entry to a specified note. The **t_getnote** directive allows the user to obtain the note contained in any one of the sixteen notepads of a specified task.

4.3.8 Task Deletion

RTEMS provides the `t_delete` directive to allow a task to delete itself or any other task. This directive removes all **RTEMS** references to the task, frees the task's control block, removes it from resource wait queues, and deallocates its stack as well as the optional floating point context. The task's name and ID become inactive at this time, and any subsequent references to either of them is invalid. In fact, **RTEMS** may reuse the task ID for another task which is created later in the application.

Unexpired timers are canceled, however, other resources dynamically allocated by the task are **NOT** automatically returned to **RTEMS**. Therefore, before a task is deleted, all of its dynamically allocated resources should be deallocated by the user. This may be accomplished by instructing the task to delete itself rather than directly deleting the task. Other tasks may instruct a task to delete itself by sending a "delete self" message, event, or signal, or by restarting the task with special arguments which instruct the task to delete itself.

4.4 Directives

This section details the task manager's directives. A subsection is dedicated to each of this manager's directives and describes the calling sequence, related constants, usage, and status codes.

4.4.1 T_CREATE - Create a task

CALLING SEQUENCE:

`dir_status t_create(name,priority,stksize,mode,attr,&tid)`

INPUT:

<code>obj_name</code>	<code>name;</code>	<code>/* user-defined name */</code>
<code>task_pri</code>	<code>priority;</code>	<code>/* task priority */</code>
<code>unsigned32</code>	<code>stksize;</code>	<code>/* stack size in bytes */</code>
<code>unsigned32</code>	<code>mode;</code>	<code>/* initial task mode */</code>
<code>unsigned32</code>	<code>attr;</code>	<code>/* task attributes */</code>

OUTPUT:

<code>obj_id</code>	<code>*tid;</code>	<code>/* id of created task */</code>
---------------------	--------------------	---------------------------------------

DIRECTIVE STATUS CODES:

SUCCESSFUL	task created successfully
E_SIZE	stack too small
E_MEMORY	no memory for stack segment
E_PRIORITY	invalid task priority
E_NOMP	multiprocessing not configured
E_TOOMANY	too many tasks created, or too many global objects

DESCRIPTION:

This directive creates a task which resides on the local node. It allocates and initializes a TCB, a stack, and an optional floating point context area. The **mode** parameter contains values which sets the task's initial execution mode. The **FP** attribute should be specified if the created task is to use a numeric coprocessor. For performance reasons, it is recommended that tasks not using the numeric coprocessor should specify the **NOFP** attribute. If the **GLOBAL** attribute is specified, the task can be accessed from remote nodes. The task id, returned in **tid**, is used in other task related directives to access the task. When created, a task is placed in the dormant state and can only be made ready to execute using the directive **t_start**.

NOTES:

This directive will not cause the calling task to be preempted.

Valid task priorities range from a high of 1 to a low of 255.

RTEMS supports a maximum of 256 interrupt levels which are mapped onto the interrupt levels actually supported by the target processor.

The requested stack size should be at least 256 bytes. Application developers should consider the stack usage of the device drivers when calculating the stack size required for tasks which utilize the driver.

The following task attribute constants are defined by **RTEMS**:

CONSTANT	DESCRIPTION	DEFAULT
NOFP	does not use coprocessor	*
FP	uses numeric coprocessor	
LOCAL	local task	*
GLOBAL	global task	

The following task mode constants are defined by **RTEMS**:

CONSTANT	DESCRIPTION	DEFAULT
PREEMPT	enable preemption	*
NOPREEMPT	disable preemption	
NOTSLICE	disable timeslicing	*
TSLICE	enable timeslicing	
ASR	enable ASR processing	*
NOASR	disable ASR processing	
INTR(0)	enable all interrupts	*
INTR(n)	execute at interrupt level n	

Tasks should not be made global unless remote tasks must interact with them. This avoids the system overhead incurred by the creation of a global task. When a global task is created, the task's name and id must be transmitted to every node in the system for insertion in the local copy of the global object table.

The total number of global objects, including tasks, is limited by the **num_objects** field in the **Configuration Table**.

4.4.2 T_IDENT - Get ID of a task

CALLING SEQUENCE:

`dir_status t_ident(name, node, &tid)`

INPUT:

`obj_name` `name;` `/* user-defined task name */`
`unsigned32` `node;` `/* nodes to be searched */`

OUTPUT:

`obj_id` `*tid;` `/* task id returned */`

DIRECTIVE STATUS CODES:

<code>SUCCESSFUL</code>	task identified successfully
<code>E_NAME</code>	invalid task name
<code>E_NODE</code>	invalid node id

DESCRIPTION:

This directive obtains the task id associated with the task name specified in `name`. A task may obtain its own id by specifying `SELF` or its own task name in `name`. If the task name is not unique, then the task id returned will match one of the tasks with that name. However, this task id is not guaranteed to correspond to the desired task. The task id, returned in `tid`, is used in other task related directives to access the task.

NOTES:

This directive will not cause the running task to be preempted.

If `node` is `ALL_NODES`, all nodes are searched with the local node being searched first. All other nodes are searched with the lowest numbered node searched first.

If `node` is a valid node number which does not represent the local node, then only the tasks exported by the designated node are searched.

This directive does not generate activity on remote nodes. It accesses only the local copy of the global object table.

4.4.3 T_START - Start a task

CALLING SEQUENCE:

`dir_status t_start(tid, saddr, arg)`

INPUT:

<code>obj_id</code>	<code>tid;</code>	<code>/* task id</code>	<code>*/</code>
<code>task_ptr</code>	<code>saddr;</code>	<code>/* task entry point</code>	<code>*/</code>
<code>unsigned32</code>	<code>arg;</code>	<code>/* argument</code>	<code>*/</code>

OUTPUT: NONE

DIRECTIVE STATUS CODES:

<code>SUCCESSFUL</code>	task started successfully
<code>E_ID</code>	invalid task id
<code>E_STATE</code>	task not in the dormant state
<code>E_REMOTE</code>	cannot start remote task

DESCRIPTION:

This directive readies the task, specified by `tid`, for execution based on the priority and execution mode specified when the task was created. The starting address of the task is given in `saddr`. The task's starting argument is contained in `arg`. This argument can be a single value or the address of a parameter block.

NOTES:

The calling task will be preempted if its preemption mode is enabled and the task being started has a higher priority.

Any actions performed on a dormant task such as suspension or change of priority are nullified when the task is initiated via the `t_start` directive.

4.4.4 T_RESTART - Restart a task

CALLING SEQUENCE:

`dir_status t_restart(tid, arg)`

INPUT:

`obj_id tid; /* task id */`
`unsigned32 arg; /* argument */`

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	task restarted successfully
E_ID	task id invalid
E_STATE	task never started
E_REMOTE	cannot restart remote task

DESCRIPTION:

This directive resets the task specified by `tid` to begin execution at its original starting address. The task's priority and execution mode are set to the original creation values. If the task is currently blocked, **RTEMS** automatically makes the task ready. A task can be restarted from any state, except the dormant state.

The task's starting argument is contained in `arg`. This argument can be a single value or the address of a parameter block. This new argument may be used to distinguish between the initial `t_start` of the task and any ensuing calls to `t_restart` of the task. This can be beneficial in deleting a task. Instead of deleting a task using the `t_delete` directive, a task can delete another task by restarting that task, and allowing that task to release resources back to **RTEMS** and then delete itself.

NOTES:

If `tid` is **SELF**, the calling task will be restarted and will not return from this directive.

The calling task will be preempted if its preemption mode is enabled and the task being restarted has a higher priority.

The task must reside on the local node, even if the task was created with the **GLOBAL** option.

4.4.5 T_DELETE - Delete a task

CALLING SEQUENCE:

```
dir_status t_delete( tid )
```

INPUT:

```
obj_id tid;      /* task id */
```

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	task deleted successfully
E_ID	invalid task id
E_REMOTE	cannot delete remote task

DESCRIPTION:

This directive deletes a task, either the calling task or another task, as specified by tid. RTEMS stops the execution of the task and reclaims the stack memory, any allocated timers (TMCBs), and the TCB. RTEMS does not reclaim region segments, partition buffers, or semaphores.

NOTES:

A task is responsible for releasing its resources back to RTEMS before deletion. To insure proper deallocation of resources, a task should not be deleted unless it is unable to execute or does not hold any RTEMS resources. If a task holds RTEMS resources, the task should be allowed to deallocate its resources before deletion. A task can be directed to release its resources and delete itself by restarting it with a special argument or by sending it a message, an event, or a signal.

Deletion of the current task (SELF) will force RTEMS to select a another task to execute.

When a global task is deleted, the task id must be transmitted to every node in the system for deletion from the local copy of the global object table.

The task must reside on the local node, even if the task was created with the GLOBAL option.

4.4.6 T_SUSPEND - Suspend a task

CALLING SEQUENCE:

dir_status t_suspend(tid)

INPUT:

obj_id tid; /* task id */

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	task suspended successfully
E_ID	invalid task id
E_ALREADY	task already suspended

DESCRIPTION:

This directive suspends the task specified by tid from further execution by placing it in the suspended state. This state is additive to any other blocked state that the task may already be in. The task will not execute again until another task issues the t_resume directive for this task and any blocked state has been removed.

NOTES:

This directive supports local operations only.

The requesting task can suspend itself by specifying SELF as tid. In this case, the task will be suspended and a successful return code will be returned when the task is resumed.

Suspending a global task which does not reside on the local node will generate a request to the remote node to suspend the specified task.

4.4.7 T_RESUME - Resume a task

CALLING SEQUENCE:

```
dir_status t_resume( tid )
```

INPUT:

```
obj_id tid;      /* task id */
```

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	task resumed successfully
E_ID	invalid task id
E_STATE	task not suspended

DESCRIPTION:

This directive removes the task specified by **tid** from the suspended state. If the task is in the ready state after the suspension is removed, then it will be scheduled to run. If the task is still in a blocked state after the suspension is removed, then it will remain in that blocked state.

NOTES:

This directive supports local operations only.

The running task may be preempted if its preemption mode is enabled and the local task being resumed has a higher priority.

Resuming a global task which does not reside on the local node will generate a request to the remote node to resume the specified task.

4.4.8 T_SETPRI - Set task priority

CALLING SEQUENCE:

```
dir_status t_setpri( tid, priority, &ppriority )
```

INPUT:

```
obj_id    tid;           /* task id          */
task_pri  priority;      /* new priority    */
```

OUTPUT:

```
task_pri *ppriority;      /* previous priority */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	task priority set successfully
E_ID	invalid task id
E_PRIORITY	invalid task priority

DESCRIPTION:

This directive manipulates the priority of the task specified by **tid**. A **tid** of **SELF** is used to indicate the calling task. When **priority** is not equal to **CURRENT**, the specified task's previous priority is returned in **ppriority**. When **priority** is **CURRENT**, the specified task's current priority is returned in **ppriority**. Valid priorities range from a high of 1 to a low of 255.

NOTES:

The calling task may be preempted if its preemption mode is enabled and it lowers its own priority or raises another task's priority.

Setting the priority of a global task which does not reside on the local node will generate a request to the remote node to change the priority of the specified task.

4.4.9 T_MODE - Change current task's mode

CALLING SEQUENCE:

```
dir_status t_mode( mode, mask, &pmode )
```

INPUT:

```
unsigned32 mode;      /* new mode values      */
unsigned32 mask;      /* mode fields to change */
```

OUTPUT:

```
unsigned32 *pmode;    /* previous mode      */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL always successful

DESCRIPTION:

This directive manipulates the execution mode of the calling task. A task's execution mode enables and disables preemption, timeslicing, asynchronous signal processing, as well as specifying the current interrupt level. To modify an execution mode, the mode class(es) to be changed must be specified in the **mask** parameter and the desired mode(s) must be specified in the **mode** parameter.

NOTES:

The calling task will be preempted if it enables preemption and a higher priority task is ready to run.

Enabling timeslicing has no effect if preemption is enabled.

A task can obtain its current execution mode, without modifying it, by calling this directive with a **mask** value of **CURRENT**.

To temporarily disable the processing of a valid ASR, a task should call this directive with the **NOASR** indicator specified in **mode**.

The following task mode constants are defined by RTEMS:

CONSTANT	DESCRIPTION	DEFAULT
PREEMPT	enable preemption	*
NOPREEMPT	disable preemption	
NOTSLICE	disable timeslicing	*
TSlice	enable timeslicing	
ASR	enable ASR processing	*
NOASR	disable ASR processing	
INTR(0)	enable all interrupts	*
INTR(n)	execute at interrupt level n,	

The following mask constants are defined by RTEMS:

CONSTANT	DESCRIPTION
CURRENT	obtain current mode
PREEMPTMODE	select preemption mode
TSlicEMODE	select timeslicing mode
ASRMODE	select ASR processing mode
INTRMODE	select interrupt level

4.4.10 T_GETNOTE - Get task notepad entry

CALLING SEQUENCE:

```
dir_status t_getnote( tid, notepad, &note )
```

INPUT:

```
obj_id      tid;          /* task id          */
unsigned32   notepad;      /* notepad location */
```

OUTPUT:

```
unsigned32 *note;          /* note value      */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	note obtained successfully
E_ID	invalid task id
E_NUMBER	invalid notepad location

DESCRIPTION:

This directive returns the **note** contained in the **notepad** location of the task specified by **tid**.

NOTES:

This directive supports local operations only.

This directive will not cause the running task to be preempted.

If **tid** is set to **SELF**, the calling task accesses its own notepad.

The sixteen notepad locations can be accessed using the constants **NOTEPAD_0** through **NOTEPAD_15**.

Getting a note of a global task which does not reside on the local node will generate a request to the remote node to obtain the notepad entry of the specified task.

4.4.11 T_SETNOTE - Set task notepad entry

CALLING SEQUENCE:

dir_status t_setnote(tid, notepad, note)

INPUT:

obj_id	tid;	/* task id	*/
unsigned32	notepad;	/* notepad location	*/
unsigned32	note;	/* new value for note	*/

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	task's note set successfully
E_ID	invalid task id
E_NUMBER	invalid notepad location

DESCRIPTION:

This directive sets the **notepad** entry for the task specified by **tid** to the value **note**.

NOTES:


This directive supports local operations only.

If **tid** is set to **SELF**, the calling task accesses its own notepad locations.

This directive will not cause the running task to be preempted.

The sixteen notepad locations can be accessed using the constants **NOTEPAD_0** through **NOTEPAD_15**.

Setting a notepad location of a global task which does not reside on the local node will generate a request to the remote node to set the specified notepad entry.



Interrupt Manager

5.1 Introduction

Any real-time executive must provide a mechanism for quick response to externally generated interrupts to satisfy the critical time constraints of the application. The **interrupt manager** provides this mechanism for **RTEMS**. This manager issues quick interrupt response times by providing the critical ability to alter task execution by allowing a task to be preempted upon exit from an ISR. The interrupt manager includes the following directives:

Name	Directive Description
i_enter	Enter an ISR
i_return	Return from an ISR

5.2 Background

5.2.1 i80386 Interrupt Processing

Although the i80386 supports multiple privilege levels, **RTEMS** and all user software execute at privilege level 0. This decision was made by the designer's of **RTEMS** to enhance compatibility with processors which do not provide sophisticated protection facilities like those of the i80386. This decision greatly simplifies the discussion of i80386 interrupt processing, as one need only consider interrupts without privilege transitions.

Without support hardware, only three levels (enabled, disabled, and non-maskable) of interrupt priorities are supported by the i80386 microprocessor. Interrupts are enabled when the interrupt-enable flag (IF) in the extended flags (EFLAGS) register is set. Conversely, interrupt processing is inhibited when IF is cleared. During a non-maskable interrupt, all other interrupts, including other non-maskable ones, are inhibited.

The i80386 processor looks for the occurrence of interrupts only between the end of one instruction and the beginning of the next. When the "repeat" prefix is used, interrupts may occur between repetitions. Upon recognition of an interrupt, the i80386 automatically pushes an interrupt stack frame (ISF) which consists of the current extended flags register, code segment, and extended instruction pointer. The i80386 then vectors to a user-supplied interrupt service routine (ISR). The ISR performs whatever minimum action is required to service that interrupt. Also, through system calls, the ISR may communicate and influence the scheduling of one or more tasks to respond to conditions indicated by the interrupt.




Figure 5-1 i80386 Interrupt Stack Frame

5.2.2 Mapping Processor Interrupt Levels

Although RTEMS supports 256 interrupt levels, the i80386 only supports three (one of which is non-maskable). RTEMS interrupt levels 0 corresponds to interrupts enabled, while RTEMS interrupt level 1 corresponds to interrupts disabled. All other RTEMS interrupt levels are undefined and their behavior is unpredictable.

5.2.3 Disabling of Interrupts by RTEMS

During the execution of directive calls, critical sections of code may be executed. When these sections are encountered, RTEMS disables interrupts before the execution of this section and restores them to the previous level upon completion of the section. RTEMS has been optimized to insure that interrupts

are disabled for less than ten (10) microseconds on a 16.67 Mhz i80386 with zero wait states.

Non-maskable interrupts cannot be disabled, and ISRs which execute at this level **MUST NEVER** issue **RTEMS** system calls. If a directive is invoked, unpredictable results may occur due to the inability of **RTEMS** to protect its critical sections. However, non-maskable ISRs that make no system calls may safely execute.

5.2.4 Format of an ISR

The **i_enter** and **i_return** directives must be used in conjunction with one another. An ISR which invokes **i_enter** must always exit with **i_return**. Likewise, no ISR should invoke **i_return** without previously invoking **i_enter**. If any ISR invokes an **RTEMS** directive, then all ISRs must use the interrupt manager.

The following is an example of a typical interrupt service routine:

```
demo_isr
    push    EAX                ; save task's EAX

    mov     EAX,I_ENTER        ; load function code
    int     RTEMS_TRAP         ; call i_enter

    ; save other registers as necessary (A)

    ...
    ; process the interrupt
    ...

    ; restore all registers saved at point A

    mov     EAX,I_RETURN       ; load function code
    int     RTEMS_TRAP         ; call i_return

    ; i_return does not return and restores EAX
```

Because the **EAX** register must be corrupted to invoke the **i_enter** directive, the ISR must save **EAX** (and no other general registers) on the stack before invoking this directive. The **EAX** register will be restored automatically by the **i_return** directive. Any other registers, including floating point, which are modified by the ISR should be saved after returning from the **i_enter** directive and restored before invoking the **i_return** directive. **RTEMS** does not guarantee the contents of any i80386 hardware data or address register if a system call (other than **i_enter** or **i_return**) is made while servicing an interrupt. To minimize the masking of other interrupts, the ISR should perform the minimum actions

required to service the interrupt. Other non-essential actions should be handled by application tasks.

5.2.5 Directives Allowed from an ISR

Using the interrupt manager insures that RTEMS knows when a directive is being called from an ISR. The ISR may then use system calls to synchronize itself with an application task. The synchronization may involve messages, events or signals being passed by the ISR to the desired task. The following is a list of RTEMS system calls that may be made from an ISR:

- *Task Management*
 - `t_getnote, t_setnote, t_suspend, t_resume`
- *Message, Event, and Signal Management*
 - `q_send, q_urgent`
 - `ev_send`
 - `as_send`
- *Semaphore Management*
 - `sm_v`
- *Time Management*
 - `tm_get, tm_tick`
- *Dual-Ported Memory Management*
 - `dp_ext2int, dp_int2ext`
- *Interrupt Handling and Fatal Error Management*
 - `i_enter, i_return`
 - `k_fatal`
- *Multiprocessing*
 - `pkt_arrived`

5.3 Operations

5.3.1 Entering an ISR

Since the interrupt vectors are user-installed and automatically vectored to by the processor, RTEMS is not involved in the switching of control from the

currently executing task to the appropriate ISR. The `enter` directive allows RTEMS to recognize, in a processor independent fashion, when an interrupt service routine is executing. This directive must be invoked immediately upon entering an interrupt service routine. RTEMS must recognize ISRs in order to postpone dispatch processing required by directives called from an ISR.

5.3.2 Exiting an ISR

RTEMS provides the `i_return` directive to allow the ISR to indicate that an interrupt service routine is terminating. This directive guarantees that proper task scheduling and dispatching are performed at the conclusion of an ISR. A system call made by the ISR may have readied a task of higher priority than the interrupted task. Therefore, when the ISR completes, the postponed dispatch processing must be performed.

Interrupts are nested whenever an interrupt occurs during the execution of another ISR. RTEMS supports efficient interrupt nesting by allowing the nested ISRs to terminate without performing any dispatch processing. Only when the outermost ISR terminates will the postponed dispatching occur.

5.4 Directives

This section details the interrupt manager's directives. A subsection is dedicated to each of this manager's directives and describes the calling sequence, related constants, usage, and status codes.

5.4.1 I_ENTER - Enter an ISR

CALLING SEQUENCE:

`void i_enter()`

INPUT: NONE

OUTPUT: NONE

DIRECTIVE STATUS CODES: NONE

DESCRIPTION:

This directive is called upon entering an ISR to inform the RTEMS executive an ISR has been entered. This allows RTEMS to correctly process nested interrupts. All ISRs which invoke RTEMS directives must call the `i_enter` and `i_return` directives. If no ISR on the local node invokes an RTEMS directive, then `i_enter` and `i_return` need not be used. If an ISR can be interrupted by an ISR which uses directives, then it must use both the `i_enter` and `i_return` directives.

NOTES:

This directive is callable from an ISR only.

This directive is callable only from assembly language.

The EAX register must be saved on the stack before loading the EAX register with the function code needed to call this directive.

If the ISR utilizes floating point instructions, then it must save and restore the floating point coprocessor's complete context.

5.4.2 I_RETURN - Return from an ISR

CALLING SEQUENCE:

void i_return()

INPUT: NONE

OUTPUT: NONE

DIRECTIVE STATUS CODES: NONE

DESCRIPTION:

This directive is called to exit an ISR. Since the ISR may have called a directive that caused another task with a higher priority to be ready to execute, the **i_return** directive allows **RTEMS** to dispatch the highest priority task that is ready to execute. If an ISR was entered by calling the **i_enter** directive, then it must exit with the **i_return** directive.

NOTES:

This directive is callable from an ISR only.


This directive is callable only from assembly language.

The task which was executing before the initial interrupt may be preempted if a higher priority task becomes ready to execute during the interrupt.

This directive **DOES NOT RETURN** to the caller.

An ISR must restore all registers saved after returning from the **i_enter** directive. The copy of **EAX** saved before invoking the **i_enter** directive will automatically be restored by the **i_return** directive.

If the ISR utilizes floating point instructions, then it **MUST** save and restore the floating point coprocessor's complete context.



Time Manager

6.1 Introduction

The **time manager** provides support for both clock and timer facilities. The directives provided by the time manager are:

Name	Directive Description
tm_set	Set system date and time
tm_get	Get system date and time
tm_wkafter	Wake up after interval
tm_wkwhen	Wake up when specified
tm_evafter	Send event set after interval
tm_evwhen	Send event set when specified
tm_evevery	Send periodic event set
tm_cancel	Cancel timer event
tm_tick	Announce a clock tick

6.2 Background

6.2.1 Required Support

For the features provided by the time manager to be utilized, periodic timer interrupts are required. Therefore, a real-time clock or some kind of hardware timer is necessary to create the timer interrupts. The `tm_tick` directive is normally called by the timer ISR to announce to RTEMs that a system clock tick has transpired. Elapsed time is measured in ticks. A tick is defined to be an integral number of milliseconds which is specified by the user in the Configuration Table.

6.2.2 Time and Date Data Structure

The clock facilities of the time manager operate upon a calendar time. These directives utilize the following date and time structure:

```
struct time_info {
    unsigned16 year;    /* year A.D. ; greater than 1987 */
    unsigned8  month;   /* month, 1 - 12 */
    unsigned8  day;     /* day, 1 - 31 */
    unsigned16 hour;    /* hour, 0 - 23 */
    unsigned8  minute;  /* minute, 0 - 59 */
    unsigned8  second;  /* second, 0 - 59 */
    unsigned32 ticks;   /* elapsed ticks between seconds */
};
```

6.2.3 Timer Types

The following types of timers are maintained by the time manager:

- *sleep timers*
- *event timers*
- *timeouts*

A **sleep timer** allows a task to delay for a given interval or up until a given time, and then wake and continue execution. This type of timer is created automatically by the `tm_wkafter` and `tm_wkwhen` directives and, as a result, does not have an RTEMs ID. Once activated, a sleep timer cannot be explicitly canceled. Each task may activate one and only one sleep timer at a time.

An **event timer** allows a task to send an event set to itself either after a given interval or at a given time. This type of timer is created automatically by the `tm_evafter`, `tm_evevery`, and `tm_evwhen` directives. All event timers are

assigned a unique **RTEMS** ID which can be used to cancel the timer. A task can have multiple event timers active simultaneously. The task must use the **ev_receive** directive to determine if the events have been posted.

Timeouts are a special type of timer automatically created when the timeout option is used on the **q_receive**, **ev_receive**, **sm_p** and **rn_getseg** directives. Each task may have one and only one timeout active at a time. When a timeout expires, it unblocks the task with a timeout status code.

Timers affect only the calling task, either by putting it to sleep or sending it an event set. For any particular task, multiple event timers can be used in combination with a single timeout or sleep timer. Under no circumstances, can a task have both a sleep timer and a timeout active simultaneously.

A **Timer Control Block (TMCB)** is allocated as part of creating a timer. The TMCB is used by **RTEMS** to manage the timer. TMCBs are automatically freed when the timer expires.

6.2.4 Timeslicing

Timeslicing is a task scheduling discipline in which tasks of equal priority are executed for a specific period of time before control of the CPU is passed to another task. It is also sometimes referred to as the **automatic round-robin** scheduling algorithm. The length of time allocated to each task is known as the **quantum** or **timeslice**.

The system's timeslice is defined as an integral number of ticks, and is specified in the **Configuration Table**. The timeslice is defined for the entire system of tasks, but timeslicing is enabled and disabled on a per task basis.

The **tm_tick** directive implements timeslicing by decrementing the running task's time-remaining counter when both timeslicing and preemption are enabled. If the task's timeslice has expired, then that task will be preempted if there exists a ready task of equal priority.

6.3 Operations

6.3.1 Announcing a Tick

RTEMS provides the **tm_tick** directive which is called from the user's real-time clock ISR to inform **RTEMS** that a tick has elapsed. The tick frequency value, defined in milliseconds, is a configuration parameter found in **RTEMS's Configuration Table**. **RTEMS** divides one-thousand milliseconds (one second)

by the number of milliseconds per tick to determine the number of calls to the **tm_tick** directive per second. The frequency of **tm_tick** calls determines the resolution (granularity) for all time dependent **RTEMS** actions. For example, calling **tm_tick** ten times per second yields a higher resolution than calling **tm_tick** two times per second. The **tm_tick** directive is responsible for maintaining both calendar time and the dynamic set of timers.

6.3.2 Setting and Obtaining the Time

The **tm_set** directive allows a task or an ISR to set the date and time maintained by **RTEMS**. Calendar time operations will return an error code if invoked before the date and time have been set. The **tm_get** directive allows a task or an ISR to obtain the current date and time.

6.3.3 Using a Sleep Timer

The **tm_wkafter** directive creates a sleep timer which allows a task to go to sleep for a specified interval. The task is blocked until the delay interval has elapsed, at which time the task is unblocked. A task calling the **tm_wkafter** directive with a delay interval of **YIELD** ticks will yield the processor to any other ready task of equal or greater priority and remain ready to execute.

The **tm_wkwhen** directive creates a sleep timer which allows a task to go to sleep until a specified date and time. The calling task is blocked until the specified date and time has occurred, at which time the task is unblocked.

6.3.4 Using an Event Timer

The **tm_evafter** directive creates an event timer which allows a task to be sent a specified event set after a specified interval. The **tm_evwhen** directive creates an event timer which is programmed to expire at a future date and time. The **tm_every** directive is similar to the **tm_evafter** directive except that the created timer is rearmed rather than canceled at the end of the interval. This results in an event set being sent at regular intervals rather than just a single time.

All three directives return a unique timer ID generated by **RTEMS** to the calling task.

The calling task is not blocked by either the **tm_evafter**, **tm_every**, or the **tm_evwhen** directive and must use the **ev_receive** directive to obtain the event set.

6.3.5 Canceling a Timer

The `tm_cancel` directive is used to cancel an event timer. The timer's control block is returned to the TMCB free list when the event timer is canceled. Timers are automatically canceled upon expiration.

6.4 Directives

This section details the time manager's directives. A subsection is dedicated to each of this manager's directives and describes the calling sequence, related constants, usage, and status codes.

6.4.1 TM_SET - Set system date and time

CALLING SEQUENCE:

```
dir_status tm_set( &timebuf )
```

INPUT:

```
time_buffer *timebuf;      /* date/time pointer */
```

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	date and time set successfully
E_CLOCK	invalid time buffer

DESCRIPTION:

This directive sets the system date and time. The date, time, and ticks in the `timebuf` structure are all range-checked, and an error is returned if any one is out of its valid range.

NOTES:

Years before 1988 are invalid.

The system date and time are based on the configured tick rate (number of milliseconds in a tick).

Setting the time forward may cause a higher priority task, blocked waiting on a specific time, to be made ready. In this case, the calling task will be preempted after the next clock tick.

Re-initializing RTEMS causes the system date and time to be reset to an uninitialized state. Another call to `tm_set` is required to re-initialize the system date and time to application specific specifications.

6.4.2 TM_GET - Get system date and time

CALLING SEQUENCE:

```
dir_status tm_get( &timebuf )
```

INPUT: NONE

OUTPUT:

```
time_buffer *timebuf;        /* date/time pointer */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	current time obtained successfully
E_NOTDEFINED	system date and time is not set

DESCRIPTION:

This directive obtains the system date and time. If the date and time have not been set with a previous call to **tm_set**, then the **E_NOTDEFINED** status code is returned.

NOTES:

This directive is callable from an ISR.

This directive will not cause the running task to be preempted. Re-initializing **RTEMS** causes the system date and time to be reset to an uninitialized state. Another call to **tm_set** is required to re-initialize the system date and time to application specific specifications.

6.4.3 TM_WKAFTER - Wake up after interval

CALLING SEQUENCE:

```
dir_status tm_wkafter( ticks )
```

INPUT:

```
interval ticks;      /* number of ticks to wait */
```

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL always successful

DESCRIPTION:

This directive blocks the calling task for the specified number of system clock ticks. When the requested interval has elapsed, the task is made ready. The **tm_tick** directive automatically updates the delay period.

NOTES:

Setting the system date and time with the **tm_set** directive has no effect on a **tm_wkafter** blocked task.

A task may give up the processor and remain in the ready state by specifying a value of **YIELD** in ticks.

The maximum timer interval that can be specified is the maximum value which can be represented by the **unsigned32** type.

6.4.4 TM_WKWHEN - Wake up when specified

CALLING SEQUENCE:

`dir_status tm_wkwhen(timebuf)`

INPUT:

`time_buffer *timebuf; /* date/time pointer */`

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	awakened at date/time successfully
E_CLOCK	invalid time buffer
E_NOTDEFINED	system date and time is not set

DESCRIPTION:

This directive blocks a task until the date and time specified in **timebuf**. At the requested date and time, the calling task will be unblocked and made ready to execute.

NOTES:

The ticks portion of the **timebuf** structure is ignored. The timing granularity of this directive is a second.

6.4.5 TM_EVAFTER - Send event set after interval

CALLING SEQUENCE:

```
dir_status tm_evafter( ticks, event, &tmid )
```

INPUT:

```
interval    ticks;      /* ticks until event    */
event_set   event;      /* event set          */
```

OUTPUT:

```
obj_id      *tmid;      /* id assigned to timer */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	timer event set up successfully
E_TOOMANY	too many timers allocated

DESCRIPTION:

This directive sets up a timer directing **RTEMS** to send the event set, **event**, to the calling task after **ticks** system clock ticks have elapsed. The id for the created timer is returned in **tmid**.

The calling task must call **ev_receive** to receive these events and will block until the timer expires. The **tm_tick** directive automatically adjusts the delay period.

NOTES:

This directive will not cause the calling task to be preempted.

Setting the system date and time by way of the **tm_set** directive has no effect on the countdown of the timer.

The maximum timer interval that can be specified is the maximum value which can be represented by the **unsigned32** type.

A task's timers (TMCBs) are canceled and reclaimed by **RTEMS** when the task is deleted.

6.4.6 TM_EVWHEN - Send event set when specified

CALLING SEQUENCE:

```
dir_status tm_evwhen( timebuf, event, &tmid )
```

INPUT:

```
time_buffer *timebuf;    /* time/date pointer    */
event_set   event;       /* event set      */
```

OUTPUT:

```
obj_id      *tmid;       /* id assigned to timer */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	timer event set up successfully
E_CLOCK	invalid time buffer
E_NOTDEFINED	system date and time is not set
E_TOOMANY	too many timers allocated

DESCRIPTION:

This directive sets up a timer directing **RTEMS** to send the event set, **event**, to the calling task at the date and time specified in **timebuf**. The id for the created timer is returned in **tmid**. The calling task must call the **ev_receive** directive to receive these events.

NOTES:

The ticks portion of the **timebuf** structure is set to zero. The timing granularity of this directive is a second.

A task's timers (TMCBs) are canceled and reclaimed by **RTEMS** when the task is deleted.

6.4.7 TM_EVEVERY - Send periodic event set

CALLING SEQUENCE:

```
dir_status tm_every( ticks, event, &tmid )
```

INPUT:

```
interval    ticks;      /* ticks between events */
event_set   event;      /* event set          */
```

OUTPUT:

```
obj_id      *tmid;      /* id assigned to timer */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	timer event set up successfully
E_TOOMANY	too many timers allocated
E_NUMBER	interval of zero is invalid

DESCRIPTION:

This directive sets up an interval timer directing **RTEMS** to send the event set, **event** to the calling task after every occurrence of **ticks** system clock ticks. The id for the created timer is returned in **tmid**.

The calling task must call **ev_receive** to receive these events and will block until the timer expires. The **tm_tick** directive is used to determine the period between each sending of the event set.

NOTES:

The directive **tm_cancel** must be used to stop the timer from sending the event set.

This directive will not cause the calling task to be preempted.

Setting the system date and time by way of the **tm_set** directive has no effect on the created timer.

The maximum timer interval that can be specified is the maximum value which can be represented by the **unsigned32** type.

A task's timers (TMCBs) are canceled and reclaimed by **RTEMS** when the task is deleted.

6.4.8 TM_CANCEL - Cancel timer event

CALLING SEQUENCE:

```
dir_status tm_cancel( tmid )
```

INPUT:

```
obj_id  tmid;      /* timer id */
```

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	timer event canceled successfully
E_ID	invalid timer id

DESCRIPTION:

This directive cancels the timer event specified by **tmid**. This timer event was scheduled by the **tm_evafter**, the **tm_evwhen**, or the **tm_evevery** directives.

NOTES:

This directive will not cause the calling task to be preempted.

6.4.9 TM_TICK - Announce a clock tick

CALLING SEQUENCE:

`dir_status tm_tick()`

INPUT: NONE

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL always successful

DESCRIPTION:

This directive announces to **RTEMS** that a system clock tick has occurred. The directive is usually called from the timer interrupt **ISR** of the local processor. This directive maintains the system date and time, decrements timers for delayed tasks and timeouts, and implements timeslicing.

NOTES:

This directive is typically called from an **ISR**.

The `ms_tick` and `tslice` parameters in the **Configuration Table** contain the number of milliseconds per tick and number of ticks per timeslice, respectively.



Semaphore Manager

7.1 Introduction

The **semaphore manager** utilizes standard Dijkstra counting semaphores to provide synchronization and mutual exclusion capabilities. The directives provided by the semaphore manager are:

Name	Directive Description
sm_create	Create a semaphore
sm_ident	Get ID of a semaphore
sm_delete	Delete a semaphore
sm_p	Acquire a semaphore
sm_v	Release a semaphore

7.2 Background

A semaphore can be viewed as a protected variable whose value can be modified only with the **sm_create**, **sm_p**, and **sm_v** directives. **RTEMS** supports both binary and counting semaphores. A **binary semaphore** is restricted to values of zero or one, while a **counting semaphore** can assume any non-negative integer value.

A binary semaphore can be used to control access to a single resource. In particular, it can be used to enforce mutual exclusion for a critical section in user code. In this instance, the semaphore would be created with an initial count of one to indicate that no task is executing the critical section of code. Upon entry to the critical section, a task must issue the **sm_p** directive to prevent other tasks from entering the critical section. Upon exit from the critical section, the task must issue the **sm_v** directive to allow another task to execute the critical section.

A counting semaphore can be used to control access to a pool of two or more resources. For example, access to three printers could be administered by a semaphore created with an initial count of three. When a task requires access to one of the printers, it issues the **sm_p** directive to obtain access to a printer. If a printer is not currently available, the task can wait for a printer to become available or return immediately. When the task has completed printing, it should issue the **sm_v** directive to allow other tasks access to the printer.

Task synchronization may be achieved by creating a semaphore with an initial count of zero. One task waits for the arrival of another task by issuing a **sm_p** directive when it reaches a synchronization point. The other task performs a corresponding **sm_v** operation when it reaches its synchronization point, thus unblocking the pending task.

7.2.1 Building an Attribute Set

In general, an attribute set is built by a bitwise OR of the desired attributes. The set of valid attributes is provided in the description of the **sm_create** and **sm_p** directives. An attribute listed as a default is not required to appear in the attribute OR list, although it is a good programming practice to specify default attributes. If all defaults are desired, the attribute **DEFAULTS** should be specified on this call.

This example demonstrates the **attr** parameter needed to create a local semaphore with a task priority waiting queue discipline. The **attr** parameter could be **PRIORITY** or **LOCAL|PRIORITY**. The **attr** parameter can be set to **PRIORITY** because **LOCAL** is the default for all created tasks. If a similar semaphore were to be known globally, then the **attr** parameter would be **GLOBAL|PRIORITY**.

7.3 Operations

7.3.1 Creating a Semaphore

The **sm_create** directive creates a semaphore with a user-specified name as well as an initial count. At create time the method for placing waiting tasks in the semaphore's task wait queue (FIFO or task priority) is specified. RTEMS allocates a **Semaphore Control Block (SMCB)** from the SMCB free list. This data structure is used by RTEMS to manage the newly created semaphore. Also, a unique semaphore ID is generated and returned to the calling task.

7.3.2 Obtaining Semaphore IDs

When a semaphore is created, RTEMS generates a unique semaphore ID and assigns it to the created semaphore until it is deleted. The semaphore ID may be obtained by either of two methods. First, as the result of an invocation of the **sm_create** directive, the semaphore ID is stored in a user provided location. Second, the semaphore ID may be obtained later using the **sm_ident** directive. The semaphore ID is used by other semaphore manager directives to access this semaphore.

7.3.3 Acquiring a Semaphore

The **sm_p** directive is used to acquire the specified semaphore. A simplified version of the **sm_p** directive can be described as follows:

```
if semaphore's count is greater than zero
    then decrement semaphore's count
    else wait for release of semaphore
return SUCCESSFUL
```

When the semaphore cannot be immediately acquired, one of the following situations applies:

- *By default, the calling task will wait forever to acquire the semaphore.*
- *Specifying **NOWAIT** forces an immediate return with an error status code.*
- *Specifying a timeout limits the interval the task will wait before returning with an error status code.*

If the task waits to acquire the semaphore, then it is placed in the semaphore's task wait queue in either FIFO or task priority order. All tasks waiting on a semaphore are returned an error code when the message queue is deleted.

7.3.4 Releasing a Semaphore

The `sm_v` directive is used to release the specified semaphore. A simplified version of the `sm_v` directive can be described as follows:

```
if no tasks are waiting on this semaphore
    then increment semaphore's count
    else assign semaphore to a waiting task
return SUCCESSFUL
```

7.3.5 Semaphore Deletion

The `sm_delete` directive removes a semaphore from the system and frees its control block. A semaphore can be deleted by any task that knows the semaphore's ID. As a result of this directive, all tasks blocked waiting to acquire the semaphore will be readied and returned a status code which indicates that the semaphore was deleted. Any subsequent references to the semaphore's name and ID are invalid.

7.4 Directives

This section details the semaphore manager's directives. A subsection is dedicated to each of this manager's directives and describes the calling sequence, related constants, usage, and status codes.

7.4.1 SM_CREATE - Create a semaphore

CALLING SEQUENCE:

```
dir_status sm_create( name, count, attr, &smid )
```

INPUT:

```
obj_name    name;        /* user-defined name */
unsigned32  count;       /* initial count    */
unsigned32  attr;        /* attributes      */
```

OUTPUT:

```
obj_id      *smid;       /* smid assigned    */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	semaphore created successfully
E_TOOMANY	too many semaphores created
E_NOMP	multiprocessing not configured
E_TOOMANY	too many global objects

DESCRIPTION:

This directive creates a semaphore which resides on the local node. The created semaphore has the user-defined name specified in **name** and the initial count specified in **count**. For control and maintenance of the semaphore, RTEMS allocates and initializes a SMCB. The RTEMS-assigned semaphore id is returned in **smid**. This semaphore id is used with other semaphore related directives to access the semaphore.

Specifying **PRIORITY** in **attr** causes tasks waiting for a semaphore to be serviced according to task priority. When **FIFO** is selected, tasks are serviced in First In-First Out order.

NOTES:

This directive will not cause the calling task to be preempted.

The following semaphore attribute constants are defined by RTEMS:

CONSTANT	DESCRIPTION	DEFAULT
FIFO	tasks wait by FIFO	*
PRIORITY	tasks wait by priority	
NOLIMIT	unlimited queue size	*
LIMIT	limit queue size to count	
LOCAL	local semaphore	*
GLOBAL	global semaphore	

Semaphores should not be made global unless remote tasks must interact with the created semaphore. This is to avoid the system overhead incurred by the creation of a global semaphore. When a global semaphore is created, the semaphore's name and id must be transmitted to every node in the system for insertion in the local copy of the global object table.

The total number of global objects, including semaphores, is limited by the `num_gobjects` field in the **Configuration Table**.

7.4.2 SM_IDENT - Get ID of a semaphore

CALLING SEQUENCE:

```
dir_status sm_ident( name, node, &smid )
```

INPUT:

```
obj_name    name;        /* user-defined name */
unsigned32  node;        /* node(s) to search */
```

OUTPUT:

```
obj_id      *smid;        /* semaphore id      */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	semaphore identified successfully
E_NAME	semaphore name not found
E_NODE	invalid node id

DESCRIPTION:

This directive obtains the semaphore id associated with the semaphore **name**. If the semaphore name is not unique, then the semaphore id will match one of the semaphores with that name. However, this semaphore id is not guaranteed to correspond to the desired semaphore. The semaphore id is used by other semaphore related directives to access the semaphore.

NOTES:

This directive will not cause the running task to be preempted.

If **node** is **ALL_NODES**, all nodes are searched with the local node being searched first. All other nodes are searched with the lowest numbered node searched first.

If **node** is a valid node number which does not represent the local node, then only the semaphores exported by the designated node are searched.

This directive does not generate activity on remote nodes. It accesses only the local copy of the global object table.

7.4.3 SM_DELETE - Delete a semaphore

CALLING SEQUENCE:

```
dir_status sm_delete( smid )
```

INPUT:

```
obj_id  smid;      /* semaphore id */
```

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	semaphore deleted successfully
E_ID	invalid semaphore id
E_REMOTE	cannot delete remote semaphore

DESCRIPTION:

This directive deletes the semaphore specified by **smid**. Any tasks that are waiting on this semaphore are unblocked with a status code for a deleted semaphore. The SMCB for this semaphore is reclaimed by **RTEMS**.

NOTES:

The calling task will be preempted if it enabled by the task's execution mode and a higher priority local task is waiting on the deleted semaphore. The calling task will **NOT** be preempted if all of the tasks that are waiting on the semaphore are remote tasks.

The calling task does not have to be the task that created the semaphore. Any local task that knows the semaphore id can delete the semaphore.

When a global semaphore is deleted, the semaphore id must be transmitted to every node in the system for deletion from the local copy of the global object table.

The semaphore must reside on the local node, even if the semaphore was created with the **GLOBAL** option.

Proxies, used to represent remote tasks, are reclaimed when the semaphore is deleted.

7.4.4 SM_P - Acquire a semaphore

CALLING SEQUENCE:

```
dir_status sm_p( smid, options, timeout )
```

INPUT:

```
obj_id      smid;          /* semaphore id */
unsigned32  options;       /* option set  */
interval    timeout;       /* wait interval */
```

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	semaphore obtained successfully
E_UNSATISFIED	semaphore not available
E_TIMEOUT	timed out waiting for semaphore
E_DELETE	semaphore deleted while waiting
E_ID	invalid semaphore id

DESCRIPTION:

This directive acquires the semaphore specified by **smid**. The **WAIT** and **NOWAIT** options of the **options** parameter are used to specify whether the calling task wants to wait for the semaphore to become available or return immediately. For either option, if the current semaphore count is positive, then it is decremented by one and the semaphore is successfully acquired by returning immediately with a successful return code.

If the calling task chooses to return immediately and the current semaphore count is zero or negative, then a status code indicating that the semaphore is not available is returned. If the calling task chooses to wait for a semaphore and the current semaphore count is zero or negative, then it is decremented by one and the calling task is placed on the semaphore's wait queue and blocked. If the semaphore was created with the **PRIORITY** option, then the calling task is inserted into the queue according to its priority. However, if the semaphore was created with the **FIFO** option, then the calling task is placed at the rear of the wait queue.

The **timeout** parameter specifies the maximum interval the calling task is willing to be blocked waiting for the semaphore. If it is set to **NOTIMEOUT**, then the calling task will wait forever.

NOTES:

The following semaphore acquisition option constants are defined by RTEMS:

CONSTANT	DESCRIPTION	DEFAULT
WAIT	wait for semaphore	*
NOWAIT	do NOT wait for semaphore	

Attempting to obtain a global semaphore which does not reside on the local node will generate a request to the remote node to access the semaphore. If the semaphore is not available and **NOWAIT** was not specified, then the task must be blocked until the semaphore is released. A proxy is allocated on the remote node to represent the task until the semaphore is released.

7.4.5 SM_V - Release a semaphore

CALLING SEQUENCE:

`dir_status sm_v(smid)`

INPUT:

`obj_id smid; /* semaphore id */`

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	semaphore released successfully
E_ID	invalid semaphore id

DESCRIPTION:

This directive releases the semaphore specified by **smid**. The semaphore count is incremented by one. If the count is zero or negative, then the first task on this semaphore's wait queue is removed and unblocked. The unblocked task may preempt the running task if the running task's preemption mode is enabled and the unblocked task has a higher priority than the running task.


NOTES:

This directive supports local operations only.

The calling task may be preempted if it causes a higher priority task to be made ready for execution.

Releasing a global semaphore which does not reside on the local node will generate a request telling the remote node to release the semaphore.

If the task to be unblocked resides on a different node from the semaphore, then the semaphore allocation is forwarded to the appropriate node, the waiting task is unblocked, and the proxy used to represent the task is reclaimed.



Message Manager

8.1 Introduction

The **message manager** provides communication and synchronization capabilities using **RTEMS** message queues. The directives provided by the message manager are:

Name	Directive Description
q_create	Create a queue
q_ident	Get ID of a queue
q_delete	Delete a queue
q_send	Put message at rear of a queue
q_urgent	Put message at front of a queue
q_broadcast	Broadcast N messages to a queue
q_receive	Receive message from a queue
q_flush	Flush all messages on a queue

8.2 Background

8.2.1 Messages

A message is a fixed length buffer where information can be stored to support communication. A message has a length of sixteen bytes. The information stored in a message is user-defined and can be actual data, pointer(s), or empty.

8.2.2 Message Queues

A message queue permits the passing of messages among tasks and ISRs. Message queues can contain a variable number of messages. Normally messages are sent to and received from the queue in FIFO order using the `q_send` directive. However, the `q_urgent` directive can be used to place messages at the head of a queue in LIFO order.

Synchronization can be realized because a task can wait for a message to arrive at a queue. Also, a task may poll a queue for the arrival of a message.

8.2.3 Building an Attribute Set

In general, a attribute set is built by a bitwise OR of the desired attributes. The set of valid attributes is provided in the description of the `q_create` and `q_receive` directives. An attribute listed as a default is not required to appear in the attribute OR list, although it is a good programming practice to specify default attributes. If all defaults are desired, the attribute `DEFAULTS` should be specified on this call.

This example demonstrates the `attr` parameter needed to create a local message queue with a task priority waiting queue discipline. The `attr` parameter could be `PRIORITY` or `LOCAL|PRIORITY`. The `attr` parameter can be set to `PRIORITY` because `LOCAL` is the default for all created message queues. If a similar message queue were to be known globally, then the `attr` parameter would be `GLOBAL|PRIORITY`.

8.3 Operations

8.3.1 Creating a Message Queue

The `q_create` directive creates a message queue with the user-defined name. Optionally, a limit can be placed on the number of messages allowed to be in the message queue at one time. The user may select FIFO or task priority as the

method for placing waiting tasks in the task wait queue. **RTEMS** allocates a Queue Control Block (QCB) from the QCB free list to maintain the newly created queue. **RTEMS** also generates a message queue ID which is returned to the calling task.

8.3.2 Obtaining Message Queue IDs

When a message queue is created, **RTEMS** generates a unique message queue ID. The message queue ID may be obtained by either of two methods. First, as the result of an invocation of the **q_create** directive, the queue ID is stored in a user provided location. Second, the queue ID may be obtained later using the **q_ident** directive. The queue ID is used by other message manager directives to access this message queue.

8.3.3 Receiving a Message

The **q_receive** directive attempts to retrieve a message from the specified message queue. If at least one message is in the queue, then the message is removed from the queue, copied to the caller's message buffer, and returned immediately. When messages are unavailable, one of the following situations applies:

- *By default, the calling task will wait forever for the message to arrive.*
- *Specifying the **NOWAIT** option forces an immediate return with an error status code.*
- *Specifying a timeout limits the period the task will wait before returning with an error status.*

If the task waits for a message, then it is placed in the message queue's task wait queue in either FIFO or task priority order. All tasks waiting on a message queue are returned an error code when the message queue is deleted.

8.3.4 Sending a Message

Messages can be sent to a queue with the **q_send** and **q_urgent** directives. These directives work identically when tasks are waiting to receive a message. A task is removed from the task waiting queue, unblocked, and the message is copied to a waiting task's message buffer.

When no tasks are waiting at the queue, **q_send** places the message at the rear of the message queue, while **q_urgent** places the message at the front of the

queue. The message is copied to a **RTEMS** message buffer and then placed in the message queue. Neither directive can successfully send a message to a full queue.

8.3.5 Broadcasting a Message

The **q_broadcast** directive sends the same message to every task waiting on the specified message queue as an atomic operation. The message is copied to each waiting task's message buffer and each task is unblocked. The number of tasks which were unblocked is returned to the caller.

8.3.6 Message Queue Deletion

The **q_delete** directive removes a message queue from the system and frees its control block. A message queue can be deleted by any task that knows the message queue's ID. As a result of this directive, all tasks blocked waiting to receive a message from the message queue will be readied and returned a status code which indicates that the message queue was deleted. Any subsequent references to the message queue's name and ID are invalid. Any messages waiting at the message queue are also deleted and deallocated.

8.4 Directives

This section details the message manager's directives. A subsection is dedicated to each of this manager's directives and describes the calling sequence, related constants, usage, and status codes.

8.4.1 Q_CREATE - Create a queue

CALLING SEQUENCE:

```
dir_status q_create( name, count, attr, &qid )
```

INPUT:

```
obj_name    name;        /* user-defined name */
unsigned32  count;       /* max message count */
unsigned32  attr;        /* queue attributes */
obj_id      *qid;        /* queue id          */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	queue created successfully
E_TOOMANY	too many queues created
E_UNSATISFIED	out of message buffers
E_NOMP	multiprocessing not configured
E_TOOMANY	too many global objects

DESCRIPTION:

This directive creates a message queue which resides on the local node with the user-defined name specified in **name**. For control and maintenance of the queue, RTEMs allocates and initializes a QCB. The RTEMs-assigned queue id, returned in **qid**, is used to access the message queue.

Specifying **PRIORITY** in **attr** causes tasks waiting for a message to be serviced according to task priority. When **FIFO** is specified, waiting tasks are serviced in First In-First Out order.

If **LIMIT** is specified in **attr**, then a limit is fixed on the maximum number of message buffers that can be contained in the queue. Buffers are dynamically allocated from the system message buffer pool as needed. The **count** parameter is disregarded if **NOLIMIT** is specified.

NOTES:

This directive will not cause the calling task to be preempted.

If the **LIMIT** option is selected and **count** has a value of zero, then the **q_send** and **q_urgent** directives will fail unless one or more tasks are already waiting on the queue.

The following message queue attribute constants are defined by RTEMS:

CONSTANT	DESCRIPTION	DEFAULT
FIFO	tasks wait by FIFO	*
PRIORITY	tasks wait by priority	
NOLIMIT	unlimited queue size	*
LIMIT	limit queue size to count	
LOCAL	local message queue	*
GLOBAL	global message queue	

Message queues should not be made global unless remote tasks must interact with the created message queue. This is to avoid the system overhead incurred by the creation of a global message queue. When a global message queue is created, the message queue's name and id must be transmitted to every node in the system for insertion in the local copy of the global object table.

The total number of global objects, including message queues, is limited by the **num_objects** field in the configuration table.

8.4.2 Q_IDENT - Get ID of a queue

CALLING SEQUENCE:

```
dir_status q_ident( name, node, &qid )
```

INPUT:

```
obj_name    name;        /* user-defined name */
unsigned32  node;        /* node(s) to search */
```

OUTPUT:

```
obj_id      *qid;        /* queue id          */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	queue identified successfully
E_NAME	queue name not found
E_NODE	invalid node id

DESCRIPTION:

This directive obtains the queue id associated with the queue name specified in **name**. If the queue name is not unique, then the queue id will match one of the queues with that name. However, this queue id is not guaranteed to correspond to the desired queue. The queue id is used with other message related directives to access the message queue.

NOTES:

This directive will not cause the running task to be preempted.

If **node** is **ALL_NODES**, all nodes are searched with the local node being searched first. All other nodes are searched with the lowest numbered node searched first.

If **node** is a valid node number which does not represent the local node, then only the message queues exported by the designated node are searched.

This directive does not generate activity on remote nodes. It accesses only the local copy of the global object table.

8.4.3 Q_DELETE - Delete a queue

CALLING SEQUENCE:

```
dir_status q_delete( qid )
```

INPUT:

```
obj_id qid;      /* queue id */
```

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	queue deleted successfully
E_ID	invalid queue id
E_REMOTE	cannot delete remote queue

DESCRIPTION:

This directive deletes the message queue specified by `qid`. Any tasks that are waiting on this queue are unblocked with an error code for a deleted queue. If no tasks are waiting, but the queue contains messages, then **RTEMS** returns these message buffers back to the system message buffer pool. The QCB for this queue is reclaimed by **RTEMS**.

NOTES:

The calling task will be preempted if its preemption mode is enabled and one or more local tasks with a higher priority than the calling task are waiting on the deleted queue. The calling task will **NOT** be preempted if the tasks that are waiting are remote tasks.

The calling task does not have to be the task that created the queue, although the task and queue must reside on the same node.

When the queue is deleted, any messages in the queue are returned to the free message buffer pool. Any information stored in those messages is lost.

When a global message queue is deleted, the message queue id must be transmitted to every node in the system for deletion from the local copy of the global object table.

Proxies, used to represent remote tasks, are reclaimed when the message queue is deleted.

8.4.4 Q_SEND - Put message at rear of a queue

CALLING SEQUENCE:

```
dir_status q_send( qid, buffer )
```

INPUT:

```
obj_id  qid;                /* queue id                */
long    (*buffer)[4];       /* message buffer pointer */
```

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	message sent successfully
E_ID	invalid queue id
E_UNSATISFIED	out of message buffers
E_TOOMANY	queue's limit has been reached

DESCRIPTION:

This directive sends the message contained in **buffer** to the queue specified by **qid**. If a task is waiting at the queue, then the message is copied to the waiting task's buffer and the task is unblocked. If no tasks are waiting at the queue, then the message is copied to a message buffer which is obtained from **RTEMS'** message buffer pool. The message buffer is then placed at the rear of the queue.

NOTES:

This directive supports local operations only.

The calling task will be preempted if it has preemption enabled and a higher priority task is unblocked as the result of this directive.

Sending a message to a global message queue which does not reside on the local node will generate a request to the remote node to post the message on the specified message queue.

If the task to be unblocked resides on a different node from the message queue, then the message is forwarded to the appropriate node, the waiting task is unblocked, and the proxy used to represent the task is reclaimed.

8.4.5 Q_URGENT - Put message at front of a queue

CALLING SEQUENCE:

```
dir_status q_urgent( qid, buffer )
```

INPUT:

```
obj_id  qid;           /* queue id           */
long    (*buffer)[4];  /* message buffer pointer */
```

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	message sent successfully
E_ID	invalid queue id
E_UNSATISFIED	out of message buffers
E_TOOMANY	queue's limit has been reached

DESCRIPTION:

This directive sends the message contained in **buffer** to the queue specified by **qid**. If a task is waiting on the queue, then the message is copied to the task's buffer and the task is unblocked. If no tasks are waiting on the queue, then the message is copied to a message buffer which is obtained from **RTEMS'** message buffer pool. The message buffer is then placed at the front of the queue.

NOTES:

This directive supports local operations only.

The calling task will be preempted if it has preemption enabled and a higher priority task is unblocked as the result of this directive.

Sending a message to a global message queue which does not reside on the local node will generate a request telling the remote node to post the message on the specified message queue.

If the task to be unblocked resides on a different node from the message queue, then the message is forwarded to the appropriate node, the waiting task is unblocked, and the proxy used to represent the task is reclaimed.

8.4.6 Q_BROADCAST - Broadcast N messages to a queue

CALLING SEQUENCE:

```
dir_status q_broadcast( qid, buffer, &count )
```

INPUT:

```
obj_id      qid;                /* queue id                */
long        (*buffer)[4];       /* message buffer pointer */
```

OUTPUT:

```
unsigned32 *count;               /* tasks made ready      */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	message broadcasted successfully
E_ID	invalid queue id

DESCRIPTION:

This directive causes all tasks that are waiting at the queue specified by **qid** to be unblocked with the message contained in **buffer**. Before a task is unblocked, the message in **buffer** is copied to that task's message buffer. The number of tasks that were unblocked is returned in **count**.

NOTES:

The calling task will be preempted if it has preemption enabled and a higher priority task is unblocked as the result of this directive.

The cost of this directive is directly related to the number of tasks waiting on the message queue, although it is more efficient than the equivalent number of invocations of **q_send**.

Broadcasting a message to a global message queue which does not reside on the local node will generate a request telling the remote node to post the message on the specified message queue.

When a task is unblocked which resides on a different node from the message queue, a copy of the message is forwarded to the appropriate node, the waiting task is unblocked, and the proxy used to represent the task is reclaimed.

8.4.7 Q_RECEIVE - Receive message from a queue

CALLING SEQUENCE:

`dir_status q_receive(qid, buffer, options, timeout)`

INPUT:

<code>obj_id</code>	<code>qid;</code>	<code>/* queue id</code>	<code>*/</code>
<code>long</code>	<code>(*buffer)[4];</code>	<code>/* message buffer pointer</code>	<code>*/</code>
<code>unsigned32</code>	<code>options;</code>	<code>/* receive options</code>	<code>*/</code>
<code>interval</code>	<code>timeout;</code>	<code>/* wait interval</code>	<code>*/</code>

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	message received successfully
E_ID	invalid queue id
E_UNSATISFIED	queue is empty
E_TIMEOUT	timed out waiting for message
E_DELETE	queue deleted while waiting

DESCRIPTION:

This directive receives a message from the message queue specified in `qid`. The **WAIT** and **NOWAIT** options of the `options` parameter allow the calling task to specify whether to wait for a message to become available or return immediately. For either option, if there is at least one message in the queue, then it is copied to `buffer` and this directive returns immediately with a successful return code.

If the calling task chooses to return immediately and the queue is empty, then a status code indicating this condition is returned. If the calling task chooses to wait at the message queue and the queue is empty, then the calling task is placed on the message wait queue and blocked. If the queue was created with the **PRIORITY** option specified, then the calling task is inserted into the wait queue according to its priority. But, if the queue was created with the **FIFO** option specified, then the calling task is placed at the rear of the wait queue.

A task choosing to wait at the queue can optionally specify a timeout value in the `timeout` parameter. The `timeout` parameter specifies the maximum interval to wait before the calling task desires to be unblocked. If it is set to **NOTIMEOUT**, then the calling task will wait forever.

NOTES:

The following message receive option constants are defined by **RTEMS**:

CONSTANT	DESCRIPTION	DEFAULT
WAIT	wait for message	*
NOWAIT	do NOT wait for message	

Receiving a message from a global message queue which does not reside on the local node will generate a request to the remote node to obtain a message from the specified message queue. If no message is available and **WAIT** was specified, then the task must be blocked until a message is posted. A proxy is allocated on the remote node to represent the task until the message is posted.

8.4.8 Q_FLUSH - Flush all messages on a queue

CALLING SEQUENCE:

```
dir_status q_flush( qid, &count )
```

INPUT:

```
obj_id      qid;          /* queue id          */
```

OUTPUT:

```
unsigned32 *count;        /* messages flushed */
```

DIRECTIVE STATUS CODES:


SUCCESSFUL	message received successfully
E_ID	invalid queue id

DESCRIPTION:

This directive removes all pending messages from the specified queue **qid**. The number of messages removed is returned in **count**. If no messages are present on the queue, **count** is set to zero.

NOTES:

Flushing all messages on a global message queue which does not reside on the local node will generate a request to the remote node to actually flush the specified message queue.



Event Manager

9.1 Introduction

The event manager provides a high performance method of intertask communication and synchronization. The directives provided by the event manager are:

Name	Directive Description
ev_send	Send event set to a task
ev_receive	Receive event condition

9.2 Background

9.2.1 Event Sets

An **event flag** is used by a task (or ISR) to inform another task of the occurrence of a significant situation. Thirty-two event flags are associated with each task. A collection of one or more event flags is referred to as an **event set**. The application developer should remember the following key characteristics of event operations when utilizing the event manager:

- *Events provide a simple synchronization facility.*
- *Events are aimed at tasks.*
- *Tasks can wait on more than one event simultaneously.*

- *Events are independent of one another.*
- *Events do not hold or transport data.*
- *Events are not queued. In other words, if an event is sent more than once before being received, the second and subsequent send operations have no effect.*

An event set is **posted** when it is directed (or sent) to a task. A **pending event** is an event that has been posted but not received. An **event condition** is used to specify the events which the task desires to receive and the algorithm which will be used to determine when the request is satisfied. An event condition is satisfied based upon one of two algorithms which are selected by the user. The **ANY** algorithm states that an event condition is satisfied when at least a single requested event is posted. The **ALL** algorithm states that an event condition is satisfied when every requested event is posted.

9.2.2 Building an Event Set or Condition

An event set or condition is built by a bitwise OR of the desired events. The set of valid events is **EVENT_0** through **EVENT_31**. If an event is not explicitly specified in the set or condition, then it is not present.

For example, when sending the event set consisting of **EVENT_6**, **EVENT_15**, and **EVENT_31**, the **event** parameter to the **ev_send** directive should be **EVENT_6|EVENT_15|EVENT_31**.

9.2.3 Building a Flag

In general, a flag is built by a bitwise OR of the desired options. The set of valid options is provided in the description of the **ev_receive** directive. An option listed as a default is not required to appear in the option OR list, although it is a good programming practice to specify default options. If all defaults are desired, the option **DEFAULTS** should be specified on this call.

This example demonstrates the **flag** parameter needed to poll for all events in a particular event condition to arrive. The **flag** parameter should be **ALL|NOWAIT** or **NOWAIT**. The **flag** parameter can be set to **NOWAIT** because **ALL** is the default condition for **ev_receive**.

9.3 Operations

9.3.1 Sending an Event Set

The `ev_send` directive allows a task (or an ISR) to direct an event set to a target task. Based upon the state of the target task, one of the following situations applies:

- *Target Task is Blocked Waiting for Events*
 - If the waiting task's input event condition is satisfied, then the task is made ready for execution.
 - If the waiting task's input event condition is not satisfied, then the event set is posted but left pending and the task remains blocked.
- *Target Task is Not Waiting for Events*
 - The event set is posted and left pending.

9.3.2 Receiving an Event Set

The `ev_receive` directive is used by tasks to accept a specific input event condition. The task also specifies whether the request is satisfied when all requested events are available or any single requested event is available. If the requested event condition is satisfied by pending events, then a successful return code and the satisfying event set are returned immediately. If the condition is not satisfied, then one of the following situations applies:

- *By default, the calling task will wait forever for the event condition to be satisfied.*
- *Specifying the NOWAIT option forces an immediate return with an error status code.*
- *Specifying a timeout limits the period the task will wait before returning with an error status code.*

9.3.3 Determining the Pending Event Set

A task can determine the pending event set by calling the `ev_receive` directive with a value of `CURRENT` for the input event condition. The pending events are returned to the calling task but the event set is left unaltered.

9.4 Directives

This section details the event manager's directives. A subsection is dedicated to each of this manager's directives and describes the calling sequence, related constants, usage, and status codes.

9.4.1 EV_SEND - Send event set to a task

CALLING SEQUENCE:

```
dir_status ev_send( tid, event )
```

INPUT:

```
obj_id    tid;          /* task id          */
event_set event;        /* event set to send */
```

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	event set sent successfully
E_ID	invalid task id

DESCRIPTION:

This directive sends an event set, `event`, to the task specified by `tid`. If a blocked task's input event condition is satisfied by this directive, then it will be made ready. If its input event condition is not satisfied, then the events satisfied are updated and the events not satisfied are left pending. If the task specified by `tid` is not blocked waiting for events, then the events sent are left pending.

NOTES:

This directive supports local operations only.

Specifying **SELF** for `tid` results in the event set being sent to the calling task.

Identical events sent to a task are not queued. In other words, the second, and subsequent, posting of an event to a task before it can perform an `ev_receive` has no effect.

The calling task will be preempted if it has preemption enabled and a higher priority task is unblocked as the result of this directive.

Sending an event set to a global task which does not reside on the local node will generate a request telling the remote node to send the event set to the appropriate task.

9.4.2 EV RECEIVE - Receive event condition

CALLING SEQUENCE:

```
dir_status ev_receive( eventin, options, timeout, &eventout )
```

INPUT:

```
event_set    eventin;      /* input event condition */
unsigned32   options;      /* receive options      */
interval     timeout;      /* wait interval        */
```

OUTPUT:

```
event_set    *eventout;    /* output event set     */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	event received successfully
E_UNSATISFIED	input event not satisfied (NOWAIT)
E_TIMEOUT	timed out waiting for event

DESCRIPTION:

This directive attempts to receive the event condition specified in `eventin`. If `eventin` is set to **CURRENT**, then the current pending events are returned in `eventout` and left pending. The **WAIT** and **NOWAIT** options in the `options` parameter are used to specify whether or not the task is willing to wait for the event condition to be satisfied. **EV_ANY** and **EV_ALL** are used in the `options` parameter are used to specify whether a single event or the complete event set is necessary to satisfy the event condition. The `eventout` parameter is returned to the calling task with the value that corresponds to the events in `eventin` that were satisfied.

If pending events satisfy the event condition, then `eventout` is set to the satisfied events and the pending events in the event condition are cleared. If the event condition is not satisfied and **NOWAIT** is specified, then `eventout` is set to the currently satisfied events. If the calling task chooses to wait, then it will block waiting for the event condition.

If the calling task must wait for the event condition to be satisfied, then the `timeout` parameter is used to specify the maximum interval to wait. If it is set to **NOTIMEOUT**, then the calling task will wait forever.


NOTES:

This directive only affects the events specified in **eventin**. Any pending events that do not correspond to any of the events specified in **eventin** will be left pending.

The following event receive option constants are defined by RTEMS;

CONSTANT	DESCRIPTION	DEFAULT
WAIT	task will wait for event	*
NOWAIT	task should not wait	
EV_ALL	return after all events	*
EV_ANY	return after any events	

10



Signal Manager

10.1 Introduction

The **signal manager** provides the capabilities required for asynchronous communication. The directives provided by the signal manager are:

Name	Directive Description
as_catch	Establish an ASR
as_send	Send signal set to a task
as_enter	Enter an ASR
as_return	Return from an ASR

10.2 Background

10.2.1 Definitions

The signal manager allows a task to optionally define an **asynchronous signal routine (ASR)**. An ASR is to a task what an ISR is to an application's set of tasks. When the processor is interrupted, the execution of an application is also interrupted and an ISR is given control. Likewise, when a signal is sent to a task, that task's execution path will be "interrupted" by the ASR. Sending a signal to a task has no effect on that task's current execution state.

A **signal flag** is used by a task (or ISR) to inform another task of the occurrence of a significant situation. Thirty-two signal flags are associated with each task. A collection of one or more signals is referred to as a **signal set**. A signal set is posted when it is directed (or sent) to a task. A **pending signal** is a signal that has been sent to a task with a valid ASR, but has not been processed by that task's ASR.

10.2.2 A Comparison of ASRs and ISRs

The format of an ASR is similar to that of an ISR with the following exceptions:

- *An ISR must invoke the **i_enter** directive immediately upon entry, while an ASR must invoke the **as_enter** directive.*
- *ISRs are scheduled by the processor hardware. ASRs are scheduled by RTEMS.*
- *When an ISR is invoked, an interrupt stack frame has been built by the CPU. When an ASR is invoked, a signal stack frame (SSF) has been built by RTEMS which contains the signal set and the interrupted task's mode.*




Figure 10-1 Signal Stack Frame

- *An ISR exits by calling the **i_return** directive, while an ASR exits by calling the **as_return** directive.*

10.2.3 Building a Signal Set

A signal set is built by a bitwise OR of the desired signals. The set of valid signals is **SIGNAL_0** through **SIGNAL_31**. If a signal is not explicitly specified in the set or condition, then it is not present.

This example demonstrates the **signal** parameter used when sending the signal set consisting of **SIGNAL_6**, **SIGNAL_15**, and **SIGNAL_31**. The **signal** parameter provided to the **as_send** directive should be **SIGNAL_6|SIGNAL_15|SIGNAL_31**.

10.2.4 Building a Mode

In general, a mode is built by a bitwise OR of the desired options. The set of valid mode options is the same as those allowed with the **t_create** and **t_mode** directives. A complete list of mode options is provided in the description of the **as_catch** directive. An option listed as a default is not required to appear in the option OR list, although it is a good programming practice to specify default options. If all defaults are desired, the option **DEFAULTS** should be specified on this call.

This example demonstrates the **mode** parameter used with the **as_catch** to establish an ASR which executes at interrupt level three and is non-preemptible. The **mode** should be set to **INTR(3)|NOPREEMPT** to indicate the desired processor mode and interrupt level.

10.3 Operations

10.3.1 Establishing an ASR

The **as_catch** directive establishes an ASR for the calling task. The address of the ASR and its execution mode are specified to this directive. The ASR's mode is distinct from the task's mode. For example, the task may allow preemption, while that task's ASR may have preemption disabled. Until a task calls **as_catch** the first time, its ASR is invalid, and no signal sets can be sent to the task.

A task may invalidate its ASR and discard all pending signals by calling **as_catch** with a value of **NULL_ASR** for the ASR's address. When a task's ASR is invalid, new signal sets sent to this task are discarded.

A task may disable ASR processing (**NOASR**) via the **t_mode** directive. When a task's ASR is disabled, the signals sent to it are left pending to be processed later when the ASR is enabled.

Any directive that can be called from a task can also be called from an ASR. A task is only allowed one active ASR. Thus, each call to **as_catch** replaces the previous one.

Normally, signal processing is disabled for the ASR's execution mode, but if signal processing is enabled for the ASR, the ASR must be reentrant.

10.3.2 Sending a Signal Set

The `as_send` directive allows both tasks and ISRs to send signals to a target task. The target task and a set of signals are specified to the `as_send` directive. The sending of a signal to a task has no effect on the execution state of that task. If the task is not the currently running task, then the signals are left pending and processed by the task's ASR the next time the task is dispatched to run. The ASR is executed immediately before the task is dispatched. If the currently running task sends a signal to itself or is sent a signal from an ISR, its ASR is immediately dispatched to run provided signal processing is enabled.

If an ASR with signals enabled is preempted by another task or an ISR and a new signal set is sent, then a new copy of the ASR will be invoked, nesting the preempted ASR. Upon completion of processing the new signal set, control will return to the preempted ASR. In this situation, the ASR should be reentrant.

Like events, identical signals sent to a task are not queued. In other words, sending the same signal multiple times to a task (without any intermediate signal processing occurring for the task), has the same result as sending that signal to that task once.

10.3.3 Entering an ASR

The `as_enter` directive is provided to increase the similarity between ASRs and ISRs. This directive returns an error code if the caller is not an ASR.

10.3.4 Returning from an ASR

The `as_return` directive is used to restore the mode and execution path of the interrupted task (or ASR) to the correct context. It must be invoked from assembly language and will not return if actually invoked from an ASR. If this directive is invoked from outside an ASR, then an error code will be returned.

10.3.5 Format of an ASR

Asynchronous signals were designed to provide the capability to generate software interrupts. The processing of software interrupts parallels that of hardware interrupts. As a result, the differences between the formats of ASRs and ISRs are limited to the use of `as_enter` and `as_return` to denote entry and exit instead of `i_enter` and `i_return`. For an example of ASR, please refer to the example ISR given in the **Interrupt Manager** chapter.

10.4 Directives

This section details the signal manager's directives. A subsection is dedicated to each of this manager's directives and describes the calling sequence, related constants, usage, and status codes.

10.4.1 AS_CATCH - Establish an ASR

CALLING SEQUENCE:

```
dir_status as_catch( asraddr, mode )
```

INPUT:

```
asr_ptr    asraddr;    /* ASR entry point */
unsigned32 mode;        /* mode for ASR    */
```

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL always successful

DESCRIPTION:

This directive establishes an asynchronous signal routine (ASR) for the calling task. The `asraddr` parameter specifies the entry point of the ASR. If `asraddr` is `NULL_ASR`, the ASR for the calling task is invalidated and all pending signals are cleared. Any signals sent to a task with an invalid ASR are discarded. The `mode` parameter specifies the execution mode for the ASR. This execution mode supersedes the task's execution mode while the ASR is executing.

NOTES:

This directive will not cause the calling task to be preempted.

The following task mode constants are defined by **RTEMS**:

CONSTANT	DESCRIPTION	DEFAULT
PREEMPT	enable preemption	*
NOPREEMPT	disable preemption	
NOTSLICE	disable timeslicing	*
TSLICE	enable timeslicing	
ASR	enable ASR processing	*
NOASR	disable ASR processing	
INTR(0)	enable all interrupts	*
INTR(n)	execute at interrupt level n	

10.4.2 AS_SEND - Send signal set to a task

CALLING SEQUENCE:

`dir_status as_send(tid, signal)`

INPUT:

`obj_id` `tid`; `/* task id */`
`signal_set` `signal`; `/* signal set to send */`

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	signal sent successfully
E_ID	task id invalid
E_NOTDEFINED	ASR invalid

DESCRIPTION:

This directive sends a signal set to the task specified in `tid`. The `signal` parameter contains the signal set to be sent to the task.

If a caller sends a signal set to a task with an invalid ASR, then an error code is returned to the caller. If a caller sends a signal set to a task whose ASR is valid but disabled, then the signal set will be caught and left pending for the ASR to process when it is enabled. If a caller sends a signal set to a task with an ASR that is both valid and enabled, then the signal set is caught and the ASR will execute the next time the task is dispatched to run.

NOTES:

This directive supports local operations only.

Sending a signal set to a task has no effect on that task's state. If a signal set is sent to a blocked task, then the task will remain blocked and the signals will be processed when the task becomes the running task.

Sending a signal set to a global task which does not reside on the local node will generate a request telling the remote node to send the signal set to the specified task.

10.4.3 AS_ENTER - Enter an ASR

CALLING SEQUENCE:

`dir_status as_enter()`

INPUT: NONE

OUTPUT: NONE

DIRECTIVE STATUS CODES:

E_CALLED Not called from an ASR

DESCRIPTION:

This directive is called upon entering an ASR to inform the RTEMS executive an ASR has been entered. All ASR's must call the **as_enter** and **as_return** directives. The ASR must restore all saved registers including floating point registers before the call to **as_return** is made.

NOTES:

This directive is callable from an ASR only.

This directive is callable only from assembly language.

The register save requirements for this directive are identical to those of the **i_enter** directive.

If the ASR utilizes floating point instructions, then it must save and restore the floating point coprocessor's complete context.

10.4.4 AS_RETURN - Return from an ASR

CALLING SEQUENCE:

`dir_status as_return()`

INPUT: NONE

OUTPUT: NONE

DIRECTIVE STATUS CODES:

E_CALLED Not called from an ASR

DESCRIPTION:


This directive is called by an ASR to return from an ASR. The ASR must call this directive to return to the task's correct context.

The ASR must restore all saved registers including floating point registers before the call to `as_return` is made.

NOTES:

When called from an ASR, this directive does not return. Otherwise, an error code is returned.

11



Partition Manager

11.1 Introduction

The **partition manager** provides facilities to dynamically allocate memory in fixed-size units. The directives provided by the region manager are:

Name	Directive Description
pt_create	Create a partition
pt_ident	Get ID of a partition
pt_delete	Delete a partition
pt_getbuf	Get buffer from a partition
pt_retbuf	Return buffer to a partition

11.2 Background

11.2.1 Definitions

A **partition** is a physically contiguous memory area divided into fixed-size buffers that can be dynamically allocated and deallocated.

Partitions are managed and maintained as a list of buffers. Buffers are obtained from the front of the partition's free buffer chain and returned to the rear of the same chain. When a buffer is on the free buffer chain, **RTEMS** uses eight

bytes of each buffer as the free buffer chain. When a buffer is allocated, the entire buffer is available for application use. Therefore, modifying memory that is outside of an allocated buffer could destroy the free buffer chain or the contents of an adjacent allocated buffer.

11.2.2 Building an Attribute Set

In general, an attribute set is built by a bitwise OR of the desired attributes. The set of valid attributes is provided in the description of the `pt_create` directive. An attribute listed as a default is not required to appear in the attribute OR list, although it is a good programming practice to specify default attributes. If all defaults are desired, the attribute `DEFAULTS` should be specified on this call. The `attr` parameter should be `GLOBAL` to indicate that the partition is to be known globally.

11.3 Operations

11.3.1 Creating a Partition

The `pt_create` directive creates a partition with a user-specified name. The partition's name, starting address, length and buffer size are all specified to the `pt_create` directive. RTEMS allocates a **Partition Control Block (PTCB)** from the PTCB free list. This data structure is used by RTEMS to manage the newly created partition. The number of buffers in the partition is calculated based upon the specified partition length and buffer size, and returned to the calling task along with a unique partition ID.

11.3.2 Obtaining Partition IDs

When a partition is created, RTEMS generates a unique partition ID and assigned it to the created partition until it is deleted. The partition ID may be obtained by either of two methods. First, as the result of an invocation of the `pt_create` directive, the partition ID is stored in a user provided location. Second, the partition ID may be obtained later using the `pt_ident` directive. The partition ID is used by other partition manager directives to access this partition.

11.3.3 Acquiring a Buffer

A buffer can be obtained by calling the `pt_getbuf` directive. If a buffer is available, then it is returned immediately with a successful return code.

Otherwise, an unsuccessful return code is returned immediately to the caller. Tasks cannot block to wait for a buffer to become available.

11.3.4 Releasing a Buffer

Buffers are returned to a partition's free buffer chain with the `pt_retbuf` directive. This directive returns an error status code if the returned buffer was not previously allocated from this partition.

11.3.5 Deleting a Partition

The `pt_delete` directive allows a partition to be removed and returned to RTEMS. When a partition is deleted, the PTCB for that partition is returned to the PTCB free list. A partition with buffers still allocated cannot be deleted. Any task attempting to do so will be returned an error status code.

11.4 Directives

This section details the partition manager's directives. A subsection is dedicated to each of this manager's directives and describes the calling sequence, related constants, usage, and status codes.

11.4.1 PT_CREATE - Create a partition

CALLING SEQUENCE:

`dir_status pt_create(name, paddr, length, bsize, attr, &ptid)`

INPUT:

<code>obj_name</code>	<code>name;</code>	<code>/* user-defined name</code>	<code>*/</code>
<code>unsigned8</code>	<code>*paddr;</code>	<code>/* physical start address</code>	<code>*/</code>
<code>unsigned32</code>	<code>length;</code>	<code>/* physical length in bytes</code>	<code>*/</code>
<code>unsigned32</code>	<code>bsize;</code>	<code>/* buffer size in bytes</code>	<code>*/</code>
<code>unsigned32</code>	<code>attr;</code>	<code>/* partition attributes</code>	<code>*/</code>

OUTPUT:

<code>obj_id</code>	<code>*ptid;</code>	<code>/* id assigned to partition</code>	<code>*/</code>
---------------------	---------------------	--	-----------------

DIRECTIVE STATUS CODES:

<code>SUCCESSFUL</code>	region created successfully
<code>E_TOOMANY</code>	too many partitions created
<code>E_ADDRESS</code>	address not on long-word boundary
<code>E_SIZE</code>	buffer size not a multiple of 4
<code>E_NOMP</code>	multiprocessing not configured
<code>E_TOOMANY</code>	too many global objects

DESCRIPTION:

This directive creates a partition of fixed size buffers from a physically contiguous memory space. The assigned partition id is returned in `ptid`. This partition id is used to access the partition with other partition related directives. For control and maintenance of the partition, **RTEMS** allocates a **PTCB** from the local **PTCB** free pool and initializes it.

NOTES:

This directive will not cause the calling task to be preempted.

The `paddr` and `bsize` parameters must be multiples of four (long-word aligned).

Memory from the partition is not used by **RTEMS** to store the Partition Control Block.

The following partition attribute constants are defined by **RTEMS**:

CONSTANT	DESCRIPTION	DEFAULT
LOCAL	local partition	*
GLOBAL	global partition	

The PTCB for a global partition is allocated on the local node. The memory space used for the partition must reside in shared memory.

Partitions should not be made global unless remote tasks must interact with the partition. This is to avoid the overhead incurred by the creation of a global partition. When a global partition is created, the partition's name and id must be transmitted to every node in the system for insertion in the local copy of the global object table.

The total number of global objects, including partitions, is limited by the **num_gobjects** field in the **Configuration Table**.

11.4.2 PT_IDENT - Get ID of a partition

CALLING SEQUENCE:

```
dir_status pt_ident( name, node, &ptid )
```

INPUT:

```
obj_name    name;      /* user-defined name */  
unsigned32  node;      /* node(s) to search */
```

OUTPUT:

```
obj_id      *ptid;     /* partition id      */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	partition identified successfully
E_NAME	partition name not found
E_NODE	invalid node id

DESCRIPTION:

This directive obtains the partition id associated with the partition **name**. If the partition name is not unique, then the partition id will match one of the partitions with that name. However, this partition id is not guaranteed to correspond to the desired partition. The partition id is used with other partition related directives to access the partition.

NOTES:

This directive will not cause the running task to be preempted.

If **node** is **ALL_NODES**, all nodes are searched with the local node being searched first. All other nodes are searched with the lowest numbered node searched first.

If **node** is a valid node number which does not represent the local node, then only the partitions exported by the designated node are searched.

This directive does not generate activity on remote nodes. It accesses only the local copy of the global object table.

11.4.3 PT_DELETE - Delete a partition

CALLING SEQUENCE:

`dir_status pt_delete(ptid)`

INPUT:

`obj_id ptid; /* partition id */`

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	partition deleted successfully
E_ID	invalid partition id
E_INUSE	buffers still in use
E_REMOTE	cannot delete remote partition

DESCRIPTION:

This directive deletes the partition specified by `ptid`. The partition cannot be deleted if any of its buffers are still allocated. The PTCB for the deleted partition is reclaimed by **RTEMS**.

NOTES:

This directive will not cause the calling task to be preempted.

The calling task does not have to be the task that created the partition. Any local task that knows the partition id can delete the partition.

When a global partition is deleted, the partition id must be transmitted to every node in the system for deletion from the local copy of the global object table.

The partition must reside on the local node, even if the partition was created with the **GLOBAL** option.

11.4.4 PT_GETBUF - Get buffer from a partition

CALLING SEQUENCE:

```
dir_status pt_getbuf( ptid, &bufaddr )
```

INPUT:

```
obj_id      ptid;          /* partition id   */
```

OUTPUT:

```
unsigned8  **bufaddr;      /* buffer address */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	buffer obtained successfully
E_ID	invalid partition id
E_UNSATISFIED	all buffers are allocated

DESCRIPTION:

This directive allows a buffer to be obtained from the partition specified in `ptid`. The address of the allocated buffer is returned in `bufaddr`.

NOTES:

This directive will not cause the running task to be preempted.

All buffers begin on a long-word boundary.

A task cannot wait on a buffer to become available.

Getting a buffer from a global partition which does not reside on the local node will generate a request telling the remote node to allocate a buffer from the specified partition.

11.4.5 PT_RETBUF - Return buffer to a partition

CALLING SEQUENCE:

`dir_status pt_retbuf(ptid, bufaddr)`

INPUT:

`obj_id ptid; /* partition id */`
`unsigned8 *bufaddr; /* buffer to return */`

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	buffer returned successfully
E_ID	invalid partition id
E_ADDRESS	buffer address not in partition


DESCRIPTION:

This directive returns the buffer specified by **bufaddr** to the partition specified by **ptid**.

NOTES:

This directive will not cause the running task to be preempted.

Returning a buffer to a global partition which does not reside on the local node will generate a request telling the remote node to return the buffer to the specified partition.



Region Manager

12.1 Introduction

The **region manager** provides facilities to dynamically allocate memory in variable sized units. The directives provided by the region manager are:

Name	Directive Description
rn_create	Create a region
rn_ident	Get ID of a region
rn_delete	Delete a region
rn_getseg	Get segment from a region
rn_retseg	Return segment to a region

12.2 Background

12.2.1 Definitions

A **region** makes up a physically contiguous memory space with user-defined boundaries from which variable-sized segments are dynamically allocated and deallocated. A **segment** is a variable size section of memory which is allocated in multiples of a user-defined page size. This page size is required to be a multiple of four greater than or equal to four. For example, if a request for a

350-byte segment is made in a region with 256-byte pages, then a 512-byte segment is allocated.

Regions are organized as doubly linked chains of variable sized memory blocks. Memory requests are allocated using a first-fit algorithm. If available, the requester receives the number of bytes requested (rounded up to the next page size). **RTEMS** requires some overhead from the region's memory for each segment that is allocated. Therefore, an application should only modify the memory of a segment that has been obtained from the region. The application should **NOT** modify the memory outside of any obtained segments and within the region's boundaries while the region is currently active in the system.

Upon return to the heap, the free block is coalesced with its neighbors (if free) on both sides to produce the largest possible unused block.

12.2.2 Building an Attribute Set

In general, an attribute set is built by a bitwise OR of the desired attributes. The set of valid attributes is provided in the description of the **rn_create** and **rn_getseg** directives. An attribute listed as a default is not required to appear in the attribute OR list, although it is a good programming practice to specify default attributes. If all defaults are desired, the attribute **DEFAULTS** should be specified on this call.

For example, the **attr** parameter should be **PRIORITY** to indicate that task priority should be used as the task waiting queue discipline.

12.3 Operations

12.3.1 Creating a Region

The **rn_create** directive creates a region with the user-defined name. The user may select FIFO or task priority as the method for placing waiting tasks in the task wait queue. **RTEMS** allocates a **Region Control Block (RNCB)** from the RNCB free list to maintain the newly created region. **RTEMS** also generates a unique region ID which is returned to the calling task.

It is not possible to calculate the exact number of bytes available to the user since **RTEMS** requires overhead for each segment allocated. For example, a region with one segment that is the size of the entire region has more available bytes than a region with two segments that collectively are the size of the entire region. The reason is that the region with one segment requires only the

overhead for one segment, while the other region requires the overhead for two segments.

Due to automatic coalescing, the number of segments in the region dynamically changes. Therefore, the total overhead required by RTEMS dynamically changes.

12.3.2 Obtaining Region IDs

When a region is created, RTEMS generates a unique region ID and assigns it to the created region until it is deleted. The region ID may be obtained by either of two methods. First, as the result of an invocation of the `rn_create` directive, the region ID is stored in a user provided location. Second, the region ID may be obtained later using the `rn_ident` directive. The region ID is used by other region manager directives to access this region.

12.3.3 Acquiring a Segment

The `rn_getseg` directive attempts to acquire a segment from a specified region. If the region has enough available free memory, then a segment is returned successfully to the caller. When the segment cannot be allocated, one of the following situations applies:

- *By default, the calling task will wait forever to acquire the segment.*
- *Specifying the NOWAIT option forces an immediate return with an error status code.*
- *Specifying a timeout limits the interval the task will wait before returning with an error status code.*

If the task waits for the segment, then it is placed in the region's task wait queue in either FIFO or task priority order. All tasks waiting on a region are returned an error when the message queue is deleted.

12.3.4 Releasing a Segment

When a segment is returned to a region by the `rn_retseg` directive, it is merged with its unallocated neighbors to form the largest possible segment. The first task on the wait queue is examined to determine if its segment request can now be satisfied. If so, it is given a segment and unblocked. This process is repeated until the first task's segment request cannot be satisfied.

12.3.5 Deleting a Region

A region can be removed from the system and returned to **RTEMS** with the **rn_delete** directive. When a region is deleted, its control block is returned to the RNCB free list. A region with segments still allocated is not allowed to be deleted. Any task attempting to do so will be returned an error.

12.4 Directives

This section details the region manager's directives. A subsection is dedicated to each of this manager's directives and describes the calling sequence, related constants, usage, and status codes.

12.4.1 RN_CREATE - Create a region

CALLING SEQUENCE:

```
dir_status  
rn_create( name, paddr, length, pagesize, attr, &rnid )
```

INPUT:

```
obj_name    name;           /* user-defined name */  
unsigned8   *paddr;         /* start address    */  
unsigned32   length;        /* length in bytes  */  
unsigned32   pagesize;      /* page size in bytes */  
unsigned32   attr;          /* region attributes */
```

OUTPUT:

```
obj_id      *rnid;          /* region id        */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	region created successfully
E_ADDRESS	address not on long-word boundary
E_TOOMANY	too many regions created
E_SIZE	invalid page size

DESCRIPTION:

This directive creates a region from a physically contiguous memory space. The assigned region id is returned in **rnid**. This region id is used as an argument to other region related directives to access the region.

For control and maintenance of the region, **RTEMS** allocates and initializes an **RNCB** from the **RNCB** free pool. Thus memory from the region is not used to store the **RNCB**. However, some overhead within the region is required by **RTEMS** each time a segment is constructed in the region.

Specifying **PRIORITY** in **attr** causes tasks waiting for a segment to be serviced according to task priority. Specifying **FIFO** in **attr** or selecting **DEFAULTS** will cause waiting tasks to be serviced in First In-First Out order.

The **paddr** parameter must be long-word aligned. The **pagesize** parameter must be a multiple of four greater than or equal to four.

NOTES:

This directive will not cause the calling task to be preempted.

The following region attribute constants are defined by RTEMS:

CONSTANT	DESCRIPTION	DEFAULT
FIFO	tasks wait by FIFO	*
PRIORITY	tasks wait by priority	

12.4.2 RN_IDENT - Get ID of a region

CALLING SEQUENCE:

```
dir_status rn_ident( name, &rnid )
```

INPUT:

```
obj_name  name;      /* user-defined name */
```

OUTPUT:

```
obj_id    *rnid;     /* region id          */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	region identified successfully
E_NAME	region name not found

DESCRIPTION:

This directive obtains the region id associated with the region name to be acquired. If the region name is not unique, then the region id will match one of the regions with that name. However, this region id is not guaranteed to correspond to the desired region. The region id is used to access this region in other region related directives.

NOTES:

This directive will not cause the running task to be preempted.

12.4.3 RN_DELETE - Delete a region

CALLING SEQUENCE:

dir_status rn_delete(rnid)

INPUT:

obj_id rnid; /* region id */

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	region deleted successfully
E_ID	invalid region id
E_INUSE	segments still in use

DESCRIPTION:

This directive deletes the region specified by **rnid**. The region cannot be deleted if any of its segments are still allocated. The RNCB for the deleted region is reclaimed by RTEMS.

NOTES:

This directive will not cause the calling task to be preempted.

The calling task does not have to be the task that created the region. Any local task that knows the region id can delete the region.

12.4.4 RN_GETSEG - Get segment from a region

CALLING SEQUENCE:

```
dir_status rn_getseg(rnid, size, options, timeout, &segaddr)
```

INPUT:

```
obj_id      rnid;          /* region id          */
unsigned32   size;          /* segment size in bytes */
unsigned32   options;       /* option set          */
interval     timeout;       /* wait interval        */
```

OUTPUT:

```
unsigned8 **segaddr;        /* segment address      */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	segment obtained successfully
E_ID	invalid region id
E_SIZE	request exceeds size of maximum segment
E_UNSATISFIED	segment of requested size not available
E_TIMEOUT	timed out waiting for segment

DESCRIPTION:

This directive obtains a variable size segment from the region specified by **rnid**. The address of the allocated segment is returned in **segaddr**. The **WAIT** and **NOWAIT** options of the **options** parameter are used to specify whether the calling tasks wish to wait for a segment to become available or return immediately if no segment is available. For either option, if a sufficiently sized segment is available, then the segment is successfully acquired by returning immediately with the **SUCCESSFUL** status code.

If the calling task chooses to return immediately and a segment large enough is not available, then an error code indicating this fact is returned. If the calling task chooses to wait for the segment and a segment large enough is not available, then the calling task is placed on the region's segment wait queue and blocked. If the region was created with the **PRIORITY** option, then the calling task is inserted into the wait queue according to its priority. But, if the region was created with the **FIFO** option, then the calling task is placed at the rear of the wait queue.

The **timeout** parameter specifies the maximum interval that a task is willing to wait to obtain a segment. If **timeout** is set to **NOTIMEOUT**, then the calling task will wait forever.

NOTES:

The actual length of the allocated segment may be larger than the requested size because a segment size is always a multiple of the region's pagesize.

The following segment acquisition option constants are defined by **RTEMS**:

CONSTANT	DESCRIPTION	DEFAULT
WAIT	task may wait for segment	*
NOWAIT	task may not wait	

12.4.5 RN_RETSEG - Return segment to a region

CALLING SEQUENCE:

dir_status rn_retseg(rnid, segaddr)

INPUT:

```
obj_id      rnid;          /* region id          */
unsigned8 *segaddr;        /* segment pointer */
```

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	segment returned successfully
E_ID	invalid region id
E_ADDRESS	segment address not in region

DESCRIPTION:


This directive returns the segment specified by **segaddr** to the region specified by **rnid**. The returned segment is merged with its neighbors to form the largest possible segment. The first task on the wait queue is examined to determine if its segment request can now be satisfied. If so, it is given a segment and unblocked. This process is repeated until the first task's segment request cannot be satisfied.

NOTES:

This directive will cause the calling task to be preempted if one or more local tasks are waiting for a segment and the following conditions exist:

- *a waiting task has a higher priority than the calling task*
- *the size of the segment required by the waiting task is less than or equal to the size of the segment returned.*

13



Dual-Ported Memory Manager

13.1 Introduction

The dual-ported memory manager provides a mechanism for converting addresses between internal and external representations for multiple dual-ported memory areas (DPMA). The directives provided by the dual-ported memory manager are:

Name	Directive Description
dp_create	Create a port
dp_ident	Get ID of a port
dp_delete	Delete a port
dp_2internal	Convert external to internal address
dp_2external	Convert internal to external address

13.2 Background

A dual-ported memory area (DPMA) is an contiguous block of RAM owned by a particular processor but which can be accessed by other processors in the system. The owner accesses the memory using internal addresses, while other

processors must use external addresses. **RTEMS** defines a **port** as a particular mapping of internal and external addresses.

There are two system configurations in which dual-ported memory is commonly found. The first is tightly-coupled multiprocessor computer systems where the dual-ported memory is shared between all nodes and is used for inter-node communication. The second configuration is computer systems with intelligent peripheral controllers. These controllers typically utilize the DPMA for high-performance data transfers.

13.3 Operations

13.3.1 Creating a Port

The **dp_create** directive creates a port into a DPMA with the user-defined name. The user specifies the association between internal and external representations for the port being created. **RTEMS** allocates a **Dual-Ported Memory Control Block (DPCB)** from the DPCB free list to maintain the newly created DPMA. **RTEMS** also generates a unique dual-ported memory port ID which is returned to the calling task. **RTEMS** does not initialize the dual-ported memory area or access any memory within it.

13.3.2 Obtaining Port IDs

When a port is created, **RTEMS** generates a unique port ID and assigns it to the created port until it is deleted. The port ID may be obtained by either of two methods. First, as the result of an invocation of the **dp_create** directive, the task ID is stored in a user provided location. Second, the port ID may be obtained later using the **dp_ident** directive. The port ID is used by other dual-ported memory manager directives to access this port.

13.3.3 Converting an Address

The **dp_2internal** directive is used to convert an address from external to internal representation for the specified port. The **dp_2external** directive is used to convert an address from internal to external representation for the specified port. If an attempt is made to convert an address which lies outside the specified DPMA, then the address to be converted will be returned.

13.3.4 Deleting a DPMA Port

A port can be removed from the system and returned to **RTEMS** with the **dp_delete** directive. When a port is deleted, its control block is returned to the DPCB free list.

13.4 Directives

This section details the dual-ported memory manager's directives. A subsection is dedicated to each of this manager's directives and describes the calling sequence, related constants, usage, and status codes.

13.4.1 DP_CREATE - Create a port

CALLING SEQUENCE:

```
dir_status dp_create(name, intaddr, extaddr, length, &dpid)
```

INPUT:

```
obj_name    name;           /* user-defined name          */
unsigned8    *intaddr;       /* initial internal address */
unsigned8    *extaddr;       /* initial external address */
unsigned32    length;        /* area size in bytes       */
```

OUTPUT:

```
obj_id      *dpid;          /* port id                  */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	port created successfully
E_ADDRESS	internal or external address not on long-word boundary
E_TOOMANY	too many DP memory areas created

DESCRIPTION:

This directive creates a port which resides on the local node for the specified DPMA. The assigned port id is returned in `dpid`. This port id is used as an argument to other dual-ported memory manager directives to convert addresses within this DPMA.

For control and maintenance of the port, **RTEMS** allocates and initializes an DPCB from the DPCB free pool. Thus memory from the dual-ported memory area is not used to store the DPCB.

NOTES:

The `intaddr` and `extaddr` parameters must be long-word aligned.

This directive will not cause the calling task to be preempted.

13.4.2 DP_IDENT - Get ID of a port

CALLING SEQUENCE:

```
dir_status dp_ident( name, &dpid )
```

INPUT:

```
obj_name  name;      /* user-defined name */
```

OUTPUT:

```
obj_id    *dpid;     /* port id          */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	port identified successfully
E_NAME	port name not found

DESCRIPTION:

This directive obtains the port id associated with the specified **name** to be acquired. If the port name is not unique, then the port id will match one of the DPMA's with that name. However, this port id is not guaranteed to correspond to the desired DPMA. The port id is used to access this DPMA in other dual-ported memory area related directives.

NOTES:

This directive will not cause the running task to be preempted.

13.4.3 DP_DELETE - Delete a port

CALLING SEQUENCE:

`dir_status dp_delete(dpid)`

INPUT:

`obj_id dpid; /* port id */`

OUTPUT: NONE

DIRECTIVE STATUS CODES:

SUCCESSFUL	port deleted successfully
E_ID	invalid port id

DESCRIPTION:

This directive deletes the dual-ported memory area specified by `dpid`. The DPCB for the deleted dual-ported memory area is reclaimed by **RTEMS**.

NOTES:

This directive will not cause the calling task to be preempted.

The calling task does not have to be the task that created the port. Any local task that knows the port id can delete the port.

13.4.4 DP_2INTERNAL - Convert external to internal address

CALLING SEQUENCE:

```
dir_status dp_2internal( dpid, extaddr, &intaddr )
```

INPUT:

```
obj_id      dpid;          /* port id          */  
unsigned8   *extaddr;      /* address to convert */
```

OUTPUT:

```
unsigned8 **intaddr;      /* internal address  */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL always successful

DESCRIPTION:

This directive converts a dual-ported memory address from external to internal representation for the specified port. If the given external address is invalid for the specified port, then the internal address is set to the given external address.

NOTES:

This directive is callable from an ISR.

This directive will not cause the calling task to be preempted.

13.4.5 DP_2EXTERNAL - Convert internal to external address

CALLING SEQUENCE:

```
dir_status dp_2external( dpid, intaddr, &extaddr )
```

INPUT:

```
obj_id      dpid;          /* port id          */
unsigned8   *intaddr;      /* address to convert */
```

OUTPUT:

```
unsigned8 **extaddr;      /* external address  */
```

DIRECTIVE STATUS CODES:

```
SUCCESSFUL          always successful
```

DESCRIPTION:


This directive converts a dual-ported memory address from internal to external representation so that it can be passed to owner of the DPMA represented by the specified port. If the given internal address is an invalid dual-ported address, then the external address is set to the given internal address.

NOTES:

This directive is callable from an ISR.

This directive will not cause the calling task to be preempted.

14



I/O Manager

14.1 Introduction

The input/output interface **manager** provides a well-defined mechanism for accessing device drivers and a structured methodology for organizing device drivers. The directives provided by the I/O manager are:

Name	Directive Description
de_init	Initialize a device driver
de_open	Open a device
de_close	Close a device
de_read	Read from a device
de_write	Write to a device
de_cntrl	Special device services

14.2 Background

14.2.1 Device Driver Table

Each application utilizing the **RTEMS** I/O manager must specify the address of a **Device Driver Table** in its **Configuration Table**. This table contains each

device driver's entry points. Each device driver may contain the following entry points:

- *Initialization*
- *Open*
- *Close*
- *Read*
- *Write*
- *Control*

If the device driver does not support a particular entry point, then that entry in the **Configuration Table** should be **NULL_DRIVER**. **RTEMS** will return **SUCCESSFUL** as the executive's and device driver's return code for these device driver entry points.

14.2.2 Major and Minor Device Numbers

Each call to the I/O manager must provide a device number as an argument. This device number is a thirty-two bit unsigned entity composed of a major and a minor device number. The most significant sixteen bits are the major number, and the least significant sixteen bits compose the minor number. The major number is the index of the requested driver's entry points in the **Device Driver Table**, and is used to select a specific device driver. The exact usage of the minor number is driver specific, but is commonly used to distinguish between a number of devices controlled by the same driver.

14.2.3 Device Driver Environment

Application developers, as well as device driver developers, must be aware of the following regarding the **RTEMS I/O Manager**:




Figure 14-1 Device Number Composition

- *A device driver routine executes in the context of the invoking task. Thus if the driver blocks, the invoking task blocks.*
- *The device driver is free to change the modes of the invoking task, although the driver should restore them to their original values.*
- *Device drivers can NOT be invoked from ISRs.*
- *Only local device drivers are accessible through the I/O manager.*
- *A device driver routine may invoke all other RTEMS directives, including I/O directives, on both local and global objects.*

Although the RTEMS I/O manager provides a framework for device drivers, it makes no assumptions regarding the construction or operation of a device driver.

14.2.4 Device Driver Interface

When an application invokes an I/O manager directive, RTEMS determines which device driver entry point must be invoked. The information passed by the application to RTEMS is then passed to the correct device driver entry point. RTEMS will invoke each device driver entry point with the following C calling sequence:

```
void de_entry( dev, argp, tid, rval )
unsigned32 dev;      /* device number */
unsigned8 *argp;     /* parameter block address */
obj_id tid;         /* ID of invoking task */
unsigned32 *rval;    /* driver's status area */
```

The format and contents of the parameter block are device driver and entry point dependent.

It is recommended that a device driver avoid generating error codes which conflict with those used by RTEMS. A common technique used to generate driver specific error codes is to logically OR the driver's major number with an error code.

14.2.5 Device Driver Initialization

RTEMS automatically initializes all device drivers when multitasking is initiated via the `init_exec` directive. RTEMS initializes the device drivers by invoking each device driver initialization entry point with the following parameters:

dev	corresponds to the major device number for this device driver with a minor device number of zero.
argp	will point to the Configuration Table .
tid	will contain zero.

The returned **rval** will be ignored by **RTEMS**. If the driver cannot successfully initialize the device, then it should invoke the fatal error manager.

14.3 Operations

The I/O manager provides directives which enable the application program to utilize device drivers in a standard manner. There is a direct correlation between the **RTEMS** I/O manager directives and the underlying device driver entry points.

14.4 Directives

This section details the I/O manager's directives. A subsection is dedicated to each of this manager's directives and describes the calling sequence, related constants, usage, and status codes.

14.4.1 DE_INIT - Initialize a device driver

CALLING SEQUENCE:

```
dir_status de_init( dev, argp, &rval )
```

INPUT:

```
unsigned32 dev;          /* 32-bit device number      */
unsigned8  *argp;        /* address of a driver      */
                        /* specific parameter block */
```

OUTPUT:

```
unsigned32 *rval;        /* return value from driver */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	successfully initialized
E_CALLED	called from within an ISR
E_NUMBER	invalid major device number

DESCRIPTION:

This directive calls the device driver initialization routine specified in the **Device Driver Table** for this major number. This directive is automatically invoked for each device driver when multitasking is initiated via the **init_exec** directive.

A device driver initialization module is responsible for initializing all hardware and data structures associated with a device. If necessary, it can allocate memory to be used during other operations.

NOTES:

This directive may or may not cause the calling task to be preempted. This is dependent on the device driver being initialized.

14.4.2 DE_OPEN - Open a device

CALLING SEQUENCE:

```
dir_status de_open( dev, argp, &rval )
```

INPUT:

```
unsigned32 dev;          /* 32-bit device number    */
unsigned8  *argp;        /* address of a driver    */
                        /* specific parameter block */
```

OUTPUT:

```
unsigned32 *rval;        /* return value from driver */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	device successfully opened
E_CALLED	called from within an ISR
E_NUMBER	invalid major device number

DESCRIPTION:

This directive calls the device driver open routine specified in the **Device Driver Table** for this major number. The open entry point is commonly used by device drivers to provide exclusive access to a device.

NOTES:

This directive may or may not cause the calling task to be preempted. This is dependent on the device driver being invoked.

14.4.3 DE_CLOSE - Close a device

CALLING SEQUENCE:

```
dir_status de_close( dev, argp, &rval )
```

INPUT:

```
unsigned32 dev;          /* 32-bit device number      */
unsigned8  *argp;        /* address of a driver      */
                        /* specific parameter block */
```

OUTPUT:

```
unsigned32 *rval;        /* return value from driver */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	device successfully closed
E_CALLED	called from within an ISR
E_NUMBER	invalid major device number

DESCRIPTION:

This directive calls the device driver close routine specified in the **Device Driver Table** for this major number. The close entry point is commonly used by device drivers to relinquish exclusive access to a device.

NOTES:

This directive may or may not cause the calling task to be preempted. This is dependent on the device driver being invoked.

14.4.4 DE_READ - Read from a device

CALLING SEQUENCE:

```
dir_status de_read( dev, argp, &rval )
```

INPUT:

```
unsigned32 dev;          /* 32-bit device number    */
unsigned8  *argp;        /* address of a driver    */
                        /* specific parameter block */
```

OUTPUT:

```
unsigned32 *rval;        /* return value from driver */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	device successfully read
E_CALLED	called from within an ISR
E_NUMBER	invalid major device number

DESCRIPTION:

This directive calls the device driver read routine specified in the Device Driver Table for this major number. Read operations typically require a buffer address as part of the argument parameter block. The contents of this buffer will be replaced with data from the device.

NOTES:

This directive may or may not cause the calling task to be preempted. This is dependent on the device driver being invoked.

14.4.5 DE_WRITE - Write to a device

CALLING SEQUENCE:

```
dir_status de_write( dev, argp, &rval )
```

INPUT:

```
unsigned32 dev;          /* 32-bit device number      */
unsigned8  *argp;        /* address of a driver      */
                        /* specific parameter block */
```

OUTPUT:

```
unsigned32 *rval;        /* return value from driver */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	device successfully written to
E_CALLED	called from within an ISR
E_NUMBER	invalid major device number

DESCRIPTION:

This directive calls the device driver write routine specified in the **Device Driver Table** for this major number. Write operations typically require a buffer address as part of the argument parameter block. The contents of this buffer will be sent to the device.

NOTES:

This directive may or may not cause the calling task to be preempted. This is dependent on the device driver being invoked.

14.4.6 DE_CNTRL - Special device services

CALLING SEQUENCE:

```
dir_status de_cntrl( dev, argp, &rval )
```

INPUT:

```
unsigned32 dev;          /* 32-bit device number      */
unsigned8  *argp;        /* address of a driver      */
                        /* specific parameter block */
```

OUTPUT:

```
unsigned32 *rval;        /* return value from driver */
```

DIRECTIVE STATUS CODES:

SUCCESSFUL	control function was successful
E_CALLED	called from within an ISR
E_NUMBER	invalid major device number

DESCRIPTION:

This directive calls the device driver I/O control routine specified in the **Device Driver Table** for this major number. The exact functionality of the driver entry called by this directive is driver dependent. It should not be assumed that the control entries of two device drivers are compatible. For example, an RS-232 driver I/O control operation may change the baud rate of a serial line, while an I/O control operation for a floppy disk driver may cause a seek operation.

NOTES:

This directive may or may not cause the calling task to be preempted. This is dependent on the device driver being invoked.



Fatal Error Manager

15.1 Introduction

The **fatal error manager** processes all fatal or irrecoverable errors. The directive provided by the fatal error manager is:

Name	Directive Description
k_fatal	Invoke the fatal error handler

15.2 Background

The fatal error manager is called upon detection of an irrecoverable error condition by either **RTEMS** or the application software. Fatal errors can be detected from three sources:

- *the executive (RTEMS)*
- *user system code*
- *user application code*

RTEMS automatically invokes the fatal error manager upon detection of an error it considers to be fatal. Similarly, the user should invoke the fatal error manager upon detection of a fatal error.

A user-supplied fatal error handler can be specified in the **User Extension Table** to provide access to debuggers and monitors which may be present on the target hardware. If configured, the fatal error manager will invoke a user-supplied

fatal error handler. If no user handler is configured or if the user handler returns control to the fatal error manager, then the **RTEMS** default fatal error handler is invoked.

The default fatal error handler disables processor interrupts, pushes the error code and program status register on the current stack, and halts execution on the local node.




Figure 15-1 Fatal Error Stack Frame

15.3 Operations

15.4 Announcing a Fatal Error

The **k_fatal** directive is invoked when a fatal error is detected. This directive is responsible for invoking an optional user-supplied fatal error handler and/or the **RTEMS** fatal error handler. All fatal error handlers are passed an error code to describe the error detected.

Occasionally, an application requires more sophisticated fatal error processing such as passing control to a debugger. For these cases, a user-supplied fatal error handler can be specified in the **RTEMS** configuration table. The **User Extension Table** parameter **fatal** contains the address of the fatal error handler to be executed when the **k_fatal** directive is called. If the parameter is set to **NULL_EXTENSION** or if the configured fatal error handler returns to the executive, then the default handler provided by **RTEMS** is executed. This default handler will halt execution on the processor where the error occurred.

15.5 Directives

This section details the fatal error manager's directives. A subsection is dedicated to each of this manager's directives and describes the calling sequence, related constants, usage, and status codes.

15.5.1 K_FATAL - Invoke the fatal error handler

CALLING SEQUENCE:

```
void k_fatal( errcode )
```

INPUT:

```
unsigned32 errcode;      /* fatal error code */
```

OUTPUT: NONE

DIRECTIVE STATUS CODES NONE

DESCRIPTION:

This directive processes fatal errors. If the FATAL error extension is defined in the configuration table, then the user-defined error extension is called. If configured and the provided FATAL extension returns, then the **RTEMS** fault error handler is invoked. The **RTEMS** fatal error handler pushes the error code and program status register on the current stack, disables interrupts, and halts the processor. This directive can be invoked by **RTEMS** or by the user's application code including initialization tasks, other tasks, and ISRs.

NOTES:

This directive supports local operations only.

Unless the user-defined error extension takes special actions such as restarting the calling task, this directive **WILL NOT RETURN** to the caller.

The user-defined extension for this directive may wish to initiate a global shutdown.



Scheduling Concepts

16.1 Introduction

The concept of scheduling in real-time systems dictates the ability to provide immediate response to specific external events, particularly the necessity of scheduling particular tasks to run within a specified time limit after the occurrence of an event. For example, software embedded in life-support systems used to monitor hospital patients must take instant action if a change in the patient's status is detected.

The component of **RTEMS** responsible for providing this capability is appropriately called the scheduler. The scheduler's sole purpose is to allocate the all important resource of processor time to the various tasks competing for attention. The **RTEMS** scheduler allocates the processor using a priority-based, preemptive algorithm augmented to provide round-robin characteristics within individual priority groups. The goal of this algorithm is to guarantee that the task which is executing on the processor at any point in time is the one with the **highest** priority among all tasks in the **ready** state.

There are two common methods of accomplishing the mechanics of this algorithm. Both ways involve a list or chain of tasks in the ready state. One method is to randomly place tasks in the ready chain forcing the scheduler to scan the entire chain to determine which task receives the processor. The other method is to schedule the task by placing it in the proper place on the ready chain based on the designated scheduling criteria at the time it enters the ready state. Thus, when the processor is free, the first task on the ready chain is allocated the processor. **RTEMS** schedules tasks using the second method to guarantee faster response times to external events.

16.2 Scheduling Mechanisms

RTEMS provides four mechanisms which allow the user to impact the task scheduling process:

- *user-selectable task priority level*
- *task preemption control*
- *task timeslicing control*
- *manual round-robin selection*

Each of these methods provides a powerful capability to customize sets of tasks to satisfy the unique and particular requirements encountered in custom real-time applications. Although each mechanism operates independently, there is a precedence relationship which governs the effects of scheduling modifications. The evaluation order for scheduling characteristics is always priority, preemption mode, and timeslicing. When reading the descriptions of timeslicing and manual round-robin it is important to keep in mind that preemption (if enabled) of a task by higher priority tasks will occur as required, overriding the other factors presented in the description.

16.2.1 Task Priority

The most significant of these mechanisms is the ability for the user to assign a priority level to each individual task when it is created and to alter a task's priority at run-time. RTEMS provides 255 priority levels. Level 255 is the lowest priority and level 1 is the highest. When a task is added to the ready chain, it is placed behind all other tasks of the same priority. This rule provides a round-robin within priority group scheduling characteristic. This means that in a group of equal priority tasks, tasks will execute in the order they become ready or FIFO order. Even though there are ways to manipulate and adjust task priorities, the most important rule to remember is:

The RTEMS scheduler will always select the highest priority task that is ready to run when allocating the processor to a task.

16.2.2 Preemption

Another way the user can alter the basic scheduling algorithm is by manipulating the preemption bit in the mode parameter of individual tasks. If

preemption is disabled for a task, then the task will not relinquish control of the processor until it terminates, blocks, or re-enables preemption. Even tasks which become ready to run and possess higher priority levels will not be allowed to execute. Note that the preemption setting has no effect on the manner in which a task is scheduled. It only applies once a task has control of the processor.

16.2.3 Timeslicing

Timeslicing or round-robin scheduling is an additional method which can be used to alter the basic scheduling algorithm. Like preemption, timeslicing is specified on a task by task basis. If timeslicing is enabled for a task, **RTEMS** will limit the amount of time the task can execute before the processor is allocated to another task. Each tick of the real-time clock reduces the currently running task's timeslice. When the execution time equals the timeslice, **RTEMS** will dispatch another task of the same priority to execute. If there are no other tasks of the same priority ready to execute, then the current task is allocated an additional timeslice and continues to run. Remember that a higher priority task will preempt the task (unless preemption is disabled) as soon as it is ready to run, even if the task has not used up its entire timeslice.

16.2.4 Manual Round-Robin

The final mechanism for altering the **RTEMS** scheduling algorithm is called manual round-robin. Manual round-robin is invoked by using the **tm_wkafter** directive with a time interval of **YIELD**. This allows a task to give up the processor and be immediately returned to the ready chain at the end of its priority group. If no other tasks of the same priority are ready to run, then the task does not lose control of the processor.

16.2.5 Dispatching Tasks

The dispatcher is the **RTEMS** component responsible for allocating the processor to a ready task. In order to allocate the processor to one task, it must be deallocated or retrieved from the task currently using it. This involves a concept called a context switch. To perform a context switch, the dispatcher saves the context of the current task and restores the context of the task which has been allocated to the processor. Saving and restoring a task's context is the storing/loading of all the essential information about a task to enable it to continue execution without any effects of the interruption. For example, a task's register contents must be the same when it is given the processor as they were when it was taken away. All of the information that must be saved or restored for a context switch is located either in the TCB or on the task's stacks.

Tasks that utilize a numeric coprocessor and are created with the **FP** attribute require additional operations during a context switch. These additional operations are necessary to save and restore the floating point context of **FP** tasks. To avoid unnecessary save and restore operations, the state of the numeric coprocessor is only saved when an **FP** task is dispatched and that task was not the last task to utilize the coprocessor.

16.3 Task State Transitions

Tasks in an **RTEMS** system must always be in one of the five allowable task states. These states are: **executing**, **ready**, **blocked**, **dormant**, and **non-existent**.

A task occupies the **non-existent** state before a **t_create** has been issued on its behalf. A task enters the **non-existent** state from any other state in the system when it is deleted with the **t_delete** directive. While a task occupies this state it does not have a TCB or a task ID assigned to it; therefore, no other tasks in the system may reference this task.

When a task is created via the **t_create** directive it enters the **dormant** state. This state is not entered through any other means. Although the task exists in the system, it cannot actively compete for system resources. It will remain in the **dormant** state until it is started via the **t_start** directive, at which time it

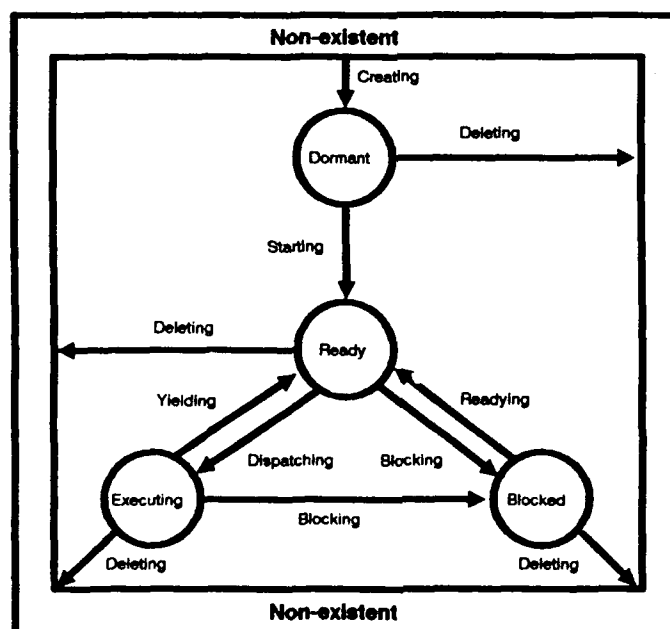


Figure 16-1 RTEMS State Transitions

enters the **ready** state. The task is now permitted to be scheduled for the processor and to compete for other system resources.

A task occupies the **blocked** state whenever it is unable to be scheduled to run. A running task may block itself or be blocked by other tasks in the system. The running task blocks itself through voluntary operations that cause the task to wait. The only way a task can block a task other than itself is with the **t_suspend** directive. A task enters the blocked state due to any of the following conditions:

- *A task issues a **t_suspend** directive which blocks either itself or another task in the system.*
- *The running task issues a **q_receive** directive with the wait option and the message queue is empty.*
- *The running task issues an **ev_receive** directive with the wait option and the currently pending events do not satisfy the request.*
- *The running task issues a **sm_p** directive with the wait option and the requested semaphore is unavailable.*
- *The running task issues a **tm_wkafter** directive which blocks the task for the given time interval. If the time interval specified is zero, the task yields the processor and remains in the ready state.*
- *The running task issues a **tm_wkwhen** directive which blocks the task until the requested date and time arrives.*
- *The running task issues a **rn_getseg** directive with the wait option and there is not an available segment large enough to satisfy the task's request.*

A blocked task may also be suspended. Therefore, both the suspension and the condition that caused the task to block, must be lifted before the task becomes ready to run.

A task occupies the **ready** state when it is able to be scheduled to run, but currently does not have control of the processor. Tasks of the same or higher priority will yield the processor by either becoming blocked, completing their timeslice, or being deleted. All tasks with the same priority will execute in **FIFO** order. A task enters the ready state due to any of the following conditions:


- *A running task issues a **t_resume** directive for a task that is suspended and the task is not blocked waiting on any resource.*

- *A running task issues a **q_send**, **q_broadcast**, or a **q_urgent** directive which posts a message to the queue on which the blocked task is waiting.*
- *A running task issues an **ev_send** directive which sends an event condition to a task which is blocked waiting on that event condition.*
- *A running task issues a **sm_v** directive which releases the semaphore on which the blocked task is waiting.*
- *A running task issues a **rn_retseg** directive which releases a segment to the region on which the blocked task is waiting and a resulting segment is large enough to satisfy the task's request.*
- *A timeout interval expires for a task which was blocked by a call to the **tm_wkafter** directive.*
- *A timeout period expires for a task which blocked by a call to the **tm_wkwhen** directive.*
- *A timeout interval expires for a task which was blocked waiting on a message, event, semaphore, or segment with a timeout specified.*
- *A running task issues a directive which deletes a message queue, a semaphore, or a region on which the blocked task is waiting.*
- *A running task issues a **t_restart** directive for the blocked task.*
- *The running task, with its preemption mode enabled, may be made **ready** by issuing any of the directives that may unblock a task with a higher priority. This directive may be issued from the running task itself or from an ISR.*

A ready task occupies the **executing** state when it has control of the CPU. A task enters the **executing** state due to any of the following conditions:

- *The task is the highest priority ready task in the system.*
- *The running task blocks and the task is next in the scheduling queue. The task may be of equal priority as in round-robin scheduling or the task may possess the highest priority of the remaining ready tasks.*

- *The running task may reenale its preemption mode and a task exists in the ready queue that has a higher priority than the running task.*
- *The running task lowers its own priority and another task is of higher priority as a result.*
- *The running task raises the priority of a task above its own and the running task is in preemption mode.*



Board Support Packages

17.1 Introduction

A board support package (BSP) is a collection of user-provided facilities which interface RTEMS with a specific hardware platform. These facilities typically include hardware initialization, Input/Output device drivers, and hardware clock management.

17.2 System Reset and Initialization

An RTEMS based application is initiated or re-initiated when the i80386 processor is reset. When the i80386 is reset, the processor performs the following actions:

- *The EAX register is set to indicate the results of the processor's power-up self test. If the self-test was not executed, the contents of this register are undefined. Otherwise, a non-zero value indicates the processor is faulty and a zero value indicates a successful self-test.*
- *The DX register holds a component identifier and revision level. DH contains 3 to indicate an i80386 component and DL contains a unique revision level indicator.*
- *Control register zero (CR0) is set such that the processor is in real mode with paging disabled. Other portions of CR0 are used to indicate the presence of a numeric coprocessor.*

- *All bits in the extended flags register (EFLAG) which are not permanently set are cleared. This inhibits all maskable interrupts.*
- *The Interrupt Descriptor Register (IDTR) is set to point at address zero.*
- *All segment registers are set to zero.*
- *The instruction pointer is set to 0x0000FFF0. The first instruction executed after a reset is actually at 0xFFFFFFF0 because the i80386 asserts the upper twelve address until the first intersegment (FAR) JMP or CALL instruction. When a JMP or CALL is executed, the upper twelve address lines are lowered and the processor begins executing in the first megabyte of memory.*

Typically, an intersegment **JMP** to the application's initialization code is placed at address 0xFFFFFFF0. Normally, the application's initialization is performed at two separate times: before the call to `init_exec` (reset application initialization) and after `init_exec` in the user's initialization tasks (local and global application initialization). The order of the startup procedure is as follows:

1. **Reset application initialization.**
2. **Call to `init_exec`**
3. **Local and global application initialization.**

The reset application initialization code is executed first when the i80386 is reset. Some of the hardware components may be initialized in this code as well as any application initialization that does not involve calls to **RTEMS** directives. This initialization code is responsible for initializing all data structures required by the i80386 in protected mode and for actually entering protected mode.

If the application requires that the **IDTR** be some value besides zero, then it should set it to the required value at this point. All tasks share the same i80386 **IDTR** value. Because interrupts are enabled automatically by **RTEMS** as part of the `init_exec` directive, the **IDTR** **MUST** be set properly before this directive is invoked to insure correct interrupt vectoring. If processor caching is to be utilized, then it should be enabled during the reset application initialization code. The reset code which is executed before the call to `init_exec` has the following requirements:

- *Must not make any RTEMS directive calls.*
- *Must set the stack segment and pointer (SS:ESP) such that a minimum stack size of 256 bytes is provided for the init_exec directive.*
- *Must leave interrupts disabled.*
- *Must zero out the executive RAM work space.*
- *Must initialize the RTEMS entry trap vector.*
- *Must initialize the Interrupt Descriptor Table.*

The `init_exec` directive does not return to the initialization code, but causes the highest priority initialization task to begin execution. Initialization tasks are used to perform both local and global application initialization which is dependent on RTEMS facilities. The user initialization task facility is typically used to create the application's set of tasks.

17.2.1 SEGMENT USAGE

RTEMS is designed to operate in the thirty-two bit flat memory model of the i80386 with paging disabled. In this mode, the i80386 automatically converts every address from a logical to a physical address each time it is used. The i80386 uses information provided in the segment registers and the Global Descriptor Table to accomplish this. RTEMS assumes the existence of following segments:

- *a single code segment at protection level zero (0) which contains all application and executive code.*
- *a single data segment at protection level zero (0) which contains all application and executive data.*

The i80386 must be placed in protected mode and the segment registers and associated selectors must be initialized before the `init_exec` directive is invoked. When each task is started or restarted, it will start with the segment selectors in use at initialization.

RTEMS does not require that logical and physical addresses are the same, although it is desirable in many applications to do so. If logical and physical addresses are different, the application may require an additional selector to access physical addresses.


Unless specified otherwise, RTEMS assumes that the DS and ES registers contain the selector for the single data segment when a directive is invoked.

This assumption is especially important when developing interrupt service routines.

17.2.2 Allowing for Interrupt Stack Usage

The i80386 does not provide for a separate stack which is used exclusively by interrupt handlers. Instead, ISRs use the stack of the task which they interrupt. Since ISRs may occur at any time, each task's stack size must take into account the worst case stack usage by any combination of nested ISRs. The following stack requirements must be accounted for:

- *Task's stack usage including RTEMS system calls*
- *Processor's interrupt stack frame*
- *RTEMS system calls requiring up to 256 bytes*
- *Registers saved on stack*
- *Application subroutine calls*



User Extensions

18.1 Introduction

RTEMS allows the application developer to augment selected features by invoking user-supplied extension routines when the following system events occur:

- *Task creation*
- *Task context switch*
- *Task initiation*
- *Task exits*
- *Task reinitiation*
- *Fatal error detection*
- *Task deletion*

These extensions are invoked as C functions with arguments that are appropriate to the system event. These functions are defined in the application's **User Extension Table** which is included as part of the **Configuration Table**. All user extensions are optional and **RTEMS** places no naming restrictions on the user.

In addition, **RTEMS** provides for a user-defined data area to be linked to each task's control block. This data area is an extension of the TCB and can be used to store additional data required by one or more of the user's extension functions. It is also possible for a user extension to utilize the notepad locations associated with each task.

The sections that follow will contain a description of each extension. Each section will contain an example of the C calling sequence for the corresponding extension. The names given for the C function and its arguments are all defined by the user. The names used in the examples were arbitrarily chosen and impose no naming conventions on the user.

18.2 TCREATE Extension

The TCREATE extension directly corresponds to the `t_create` directive. If this extension is defined in the **Configuration Table** and a task is being created, then the extension routine will automatically be invoked by RTEMS. The extension should be prototyped as follows:

```
void tcreate( curtcbl, newtcbl )
t_cb *curtcbl;
t_cb *newtcbl;
```

where `curtcbl` is the pointer to the TCB for the currently executing task, and `newtcbl` is the pointer to the TCB for the new task being created. This extension is invoked from the `t_create` directive after `newtcbl` has been completely initialized, but before it is placed on a ready TCB chain.

18.3 TSTART Extension

The TSTART extension directly corresponds to the `t_start` directive. If this extension is defined in the **Configuration Table** and a task is being started, then the extension routine will automatically be invoked by RTEMS. The extension should be prototyped as follows:

```
void tstart( curtcbl, sttcbl )
t_cb *curtcbl;
t_cb *sttcbl;
```

where `curtcbl` is the pointer to the TCB for the currently executing task, and `sttcbl` is the pointer to the TCB for the dormant task being started. This extension is invoked from the `t_start` directive after `sttcbl` has been made ready to start execution, but before it is placed on a ready TCB chain.

18.4 TRESTART Extension

The TRESTART extension directly corresponds to the `t_restart` directive. If this extension is defined in the **Configuration Table** and a task is being restarted, then the extension should be prototyped as follows:

```
void trestart( curtcbl, sttcbl )
t_cb *curtcbl;
t_cb *sttcbl;
```

where `curtcbl` is the pointer to the TCB for the currently executing task, and `sttcbl` is the pointer to the TCB for the task being restarted. This extension is invoked from the `t_restart` directive after `sttcbl` has been made ready to start execution, but before it is placed on a ready TCB chain.

18.5 TDELETE Extension

The TDELETE extension is associated with the `t_delete` directive. If this extension is defined in the **Configuration Table** and a task is being deleted, then the extension routine will automatically be invoked by RTEMS. The extension should be prototyped as follows:

```
void tdelete( curtc_b, deltc_b )
t_cb *curtc_b;
t_cb *deltc_b;
```

where `curtc_b` is the pointer to the TCB for the currently executing task, and `deltc_b` is the pointer to the TCB for the task being deleted. This extension is invoked from the `t_delete` directive after the TCB has been removed from a ready TCB chain, but before all its resources including the TCB have been returned to their respective free pools. This extension should not call any RTEMS directives if a task is deleting itself (`curtc_b` is equal to `deltc_b`).

18.6 TSWITCH Extension

The TSWITCH extension corresponds to a task context switch. If this extension is defined in the **Configuration Table** and a task context switch is in progress, then the extension routine will automatically be invoked by RTEMS. The extension should be prototyped as follows:

```
void tswitch( curtc_b, heirtc_b )
t_cb *curtc_b;
t_cb *heirtc_b;
```

where `curtc_b` is the pointer to the TCB for the task that is being swapped out, and `heirtc_b` is the pointer to the TCB for the task being swapped in. This extension is invoked from RTEMS' dispatcher routine after the `curtc_b` context has been saved, but before the `heirtc_b` context has been restored. This extension should not call any RTEMS directives.

18.7 TASKEEXITED Error Extension

The TASKEEXITED error extension is invoked when a task exits by either an implicit or explicit return statement. The address of this extension is automatically placed on each task's stack along with its argument list at task creation time. This user extension is prototyped as follows:

```
void task-xitted()
```

This extension is invoked with no arguments and should not be invoked directly. Although exiting of task is typically considered to be a fatal error, this extension

allows recovery by either restarting or deleting the exiting task. If the user does not wish to recover, then a fatal error may be reported.

If user-provided **TASKEXITTED** error extension is not provided, then the **RTEMS** default handler will be used.

18.8 FATAL Error Extension

The **FATAL** error extension is associated with the **k_fatal** directive. If this extension is defined in the **Configuration Table** and the **k_fatal** directive has been invoked, then this extension will be called. This extension should be prototyped as follows:

```
void k_fatal( errcode )
unsigned32 errcode;
```


where **errcode** is the error code passed to the **k_fatal** directive. This extension is invoked from the **k_fatal** directive.

If defined, the user's **FATAL** error extension is invoked before **RTEMS'** default fatal error routine is invoked and the processor is stopped. For example, this extension could be used to pass control to a debugger when a fatal error occurs. This extension should not call any **RTEMS** directives.

18.9 TCB Extension

The **TCB** extension is a pointer field in the **TCB** which can be set by the user to access an area of **RAM**. This allows an application to augment the **TCB** with user-defined information. For example, an application could implement task profiling by storing timing statistics in the **TCB's** extended memory area. When a task context switch is being executed, the **TSWITCH** extension could read a real-time clock to calculate how long the task being swapped out has run as well as timestamp the starting time for the task being swapped in.

If used, the extended memory area for the **TCB** should be allocated and the **TCB** extension pointer should be set at the time the task is created or started by either the **TCREATE** or **TSTART** extension. The application is responsible for managing this extended memory area for the **TCBs**. The memory may be reinitialized by the **TRESTART** extension and should be deallocated by the **TDELETE** extension when the task is deleted. Since the **TCB** extension buffers would most likely be of a fixed size, the **RTEMS** partition manager could be used to manage the application's extended memory area. The application could create a partition of fixed size **TCB** extension buffers and use the partition manager's allocation and deallocation directives to obtain and release the extension buffers.



Configuring a System

19.1 Configuration Table

The **RTEMS Configuration Table** is used to tailor an application for its specific needs. For example, the user can configure the maximum number of tasks for this application. The address of the user-defined **Configuration Table** is passed as an argument to the `init_exec` directive, which **MUST** be the first **RTEMS** directive called. The **RTEMS Configuration Table** is defined in a C structure. Each entry in the table is either two or four bytes in length. The structure is given here:

```
struct config_info {
    unsigned32    exec_ram;           /* RTEMS RAM Work Area */
    unsigned32    ram_size;           /* RTEMS Work Area size */
    unsigned16    max_tasks;          /* max number tasks */
    unsigned16    max_semaphores;     /* max number semaphores */
    unsigned16    max_timers;         /* max number timers */
    unsigned16    max_queues;         /* max number queues */
    unsigned16    max_messages;       /* max number messages */
    unsigned16    max_regions;        /* max number regions */
    unsigned16    max_partitions;     /* max number partitions */
    unsigned16    max_dpmems;         /* max dp memory areas */
    unsigned16    ms_tick;            /* ms in a tick */
    unsigned16    tslice;             /* ticks in a timeslice */
    unsigned32    num_itasks;          /* number of init tasks */
    itask_table   *Itasks_tbl;         /* init task table */
    unsigned32    num_devices;         /* number device drivers */
    driver_table  *Drv_tbl;            /* driver table */
    ext_table     *Ext_tbl;            /* extension table */
    mp_table      *Mp_tbl;             /* MP config table */
};
```

exec_ram is the starting address of the **RTEMS RAM Workspace**. This area contains items such as the various object control blocks (TCBs, QCBs, ...) and task stacks. If the address is

not aligned on a four-word boundary, then RTEMS will invoke the fatal error handler during `init_exec`.

ram_size	is the calculated size of the RTEMS RAM Workspace. The section Sizing the RTEMS RAM Workspace details how to arrive at this number.
max_tasks	is the maximum number of tasks that can be concurrently active (created) in the system including initialization tasks.
max_semaphores	is the maximum number of semaphores that can be concurrently active in the system.
max_timers	is the maximum number of timers that can be concurrently active in the system.
max_queues	is the maximum number of message queues that can be concurrently active in the system.
max_messages	is the maximum number of messages that can be allocated to the application.
max_regions	is the maximum number of regions that can be concurrently active in the system.
max_partitions	is the maximum number of partitions that can be concurrently active in the system.
max_dpmems	is the maximum number of dual-port memory areas that can be concurrently active in the system.
ms_tick	is number of milliseconds per clock tick.
tslice	is the number of clock ticks for a timeslice.
num_itasks	is the number of initialization tasks configured. At least one initialization task must be configured.
Itasks_tbl	is the address of the Initialization Task Table . This table contains the information needed to create and start each of the initialization tasks. The format of this table will be discussed below.
num_devices	is the number of device drivers for the system. There should be the same number of entries in the Device Driver

Table. If this field is zero, then the **Drv_tbl** entry should be **NULL_DRIVER_TABLE**.

Drv_tbl	is the address of the Device Driver Table . This table contains the entry points for each device driver. If the num_devices field is zero, then this entry should be NULL_DRIVER_TABLE . The format of this table will be discussed below.
Ext_tbl	is the address of the User Extension Table . This table contains the entry points for each user extension. If no user extensions are configured, then this entry should be NULL_EXT_TABLE . The format of this table will be discussed below.
Mp_tbl	is the address of the Multiprocessor Configuration Table . This table contains information needed by RTEMS only when used in a multiprocessor configuration. This field must be NULL_MP_TABLE when RTEMS is used in a single processor configuration.

19.2 Initialization Task Table

The **Initialization Task Table** is used to describe each of the user initialization tasks to the Initialization Manager. The table contains one entry for each initialization task the user wishes to create and start. The fields of the structure directly correspond to arguments to the **t_create** and **t_start** directives. The number of entries is found in the **num_itasks** entry in the **Configuration Table**. The format of each entry in the **Initialization Task Table** is defined in a C structure, and is given below:

```
struct itasks_info {
    obj_name    name;           /* task name           */
    unsigned32  stksize;        /* task stack size     */
    task_pri    priority;       /* task priority       */
    unsigned32  attributes;     /* task attributes     */
    task_ptr    entry;          /* task entry point    */
    unsigned32  mode;           /* task initial mode   */
    unsigned32  arg;            /* task argument       */
};
```

name	is the name of this initialization task.
stksize	is the size of the stack for this initialization task.
priority	is the priority of this initialization task.

attributes	is the attribute set used during creation of this initialization task.
entry	is the address of the entry point of this initialization task.
mode	is the initial execution mode of this initialization task.
arg	is the initial argument for this initialization task.

A typical declaration for an **Initialization Task Table** might appear as follows:

```
itask_table Inittasks[2] = {
    { INIT1_NAME, 1024, 1, 0,
      init1, INTR(0)|NOPREEMPT, Init1_arg },
    { INIT2_NAME, 1024, 1024, 1, 0,
      init2, INTR(0)|NOPREEMPT, Init2_arg }
};
```

19.3 Driver Address Table

The **Device Driver Table** is used to inform the I/O Manager of the set of entry points for each device driver configured in the system. The table contains one entry for each device driver required by the application. The number of entries is defined in the **num_devices** entry in the **Configuration Table**. The format of each entry in the **Device Driver Table** is defined in a C structure, and is given below:

```
struct driver_info {
    proc_ptr  init;          /* initialization procedure */
    proc_ptr  open;          /* open request procedure */
    proc_ptr  close;         /* close request procedure */
    proc_ptr  read;          /* read request procedure */
    proc_ptr  write;         /* write request procedure */
    proc_ptr  cntrl;         /* special request procedure */
    unsigned32 reserved1;    /* reserved for RTEMS use */
    unsigned32 reserved2;    /* reserved for RTEMS use */
};
```

init	is the address of the entry point called by de_init to initialize a device driver and its associated devices.
open	is the address of the entry point called by de_open .
close	is the address of the entry point called by de_close .
read	is the address of the entry point called by de_read .
write	is the address of the entry point called by de_write .

cntrl is the address of the entry point called by **de_cntrl**.

reserved1 is reserved for **RTEMS** use and should be set to **RESERVED**.

reserved2 is reserved for **RTEMS** use and should be set to **RESERVED**.

Driver entry points configured as **NULL_DRIVER** will always return a status code of **SUCCESSFUL**. No user code will be executed in this situation.

A typical declaration for a Device Driver Table might appear as follows:

```
driver_table Driver_table[2] = {
    { tty_open, tty_open, tty_close, tty_read,
      tty_write, tty_cntrl, RESERVED, RESERVED },
    { lp_open, lp_open, lp_close, NULL_DRIVER
      lp_write, lp_cntrl, RESERVED, RESERVED }
};
```

More information regarding the construction and operation of device drivers is provided in the **I/O Manager** chapter.

19.4 User Extensions Table

The **User Extensions Table** is used to inform **RTEMS** of each of the optional user-supplied extensions. This table contains one entry for each possible extension. The entries are called at critical times in the life of a task. The format of each entry in the **User Extensions Table** is defined in a C structure, and is given below:

```
struct ext_info {
    proc_ptr tcreate; /* tcreate user extension */
    proc_ptr tstart; /* tstart user extension */
    proc_ptr trestart; /* trestart user extension */
    proc_ptr tdelete; /* tdelete user extension */
    proc_ptr tswitch; /* tswitch user extension */
    proc_ptr taskexitted; /* task exit handler */
    proc_ptr fatal; /* fatal error handler */
};
```

tcreate is the address of the user-supplied subroutine for the **TCREATE** extension. If this extension for task creation is defined, it is called from the **t_create** directive. A value of **NULL_EXTENSION** indicates that no extension is provided.

tstart	is the address of the user-supplied subroutine for the TSTART extension. If this extension for task initiation is defined, it is called from the t_start directive. A value of NULL_EXTENSION indicates that no extension is provided.
trestart	is the address of the user-supplied subroutine for the TRESTART extension. If this extension for task re-initiation is defined, it is called from the t_restart directive. A value of NULL_EXTENSION indicates that no extension is provided.
tdelete	is the address of the user-supplied subroutine for the TDELETE extension. If this RTEMS extension for task deletion is defined, it is called from the t_delete directive. A value of NULL_EXTENSION indicates that no extension is provided.
tswitch	is the address of the user-supplied subroutine for the task context switch extension. This subroutine is called from RTEMS' dispatcher after the current task has been swapped out but before the new task has been swapped in. A value of NULL_EXTENSION indicates that no extension is provided. As this routine is invoked after saving the current task's context and before restoring the heir task's context, it is not necessary for this routine to save and restore any registers.
taskexitted	is the address of the user-supplied subroutine which is invoked when a task exits. This procedure is responsible for some action which will allow the system to continue execution (i.e. delete or restart the task) or to terminate with a fatal error. If this field is set to NULL_EXTENSION , the default RTEMS taskexitted handler will be invoked.
fatal	is the address of the user-supplied subroutine for the FATAL extension. This RTEMS extension of fatal error handling is called from the k_fatal directive. If the user's fatal error handler returns or if this entry is NULL_EXTENSION then the default RTEMS fatal error handler will be executed.

A typical declaration for a **User Extension Table** which defines the TCREATE, TDELETE, TSWITCH, and FATAL extension might appear as follows:


```

ext_table User_extensions = {
    tcreate_ext, NULL_EXTENSION, NULL_EXTENSION,
    tdelete_ext, tswitch_ext,
    NULL_EXTENSION, fatal_ext
};

```

More information regarding the user extensions is provided in the **User Extensions** chapter.

19.5 Multiprocessor Configuration Table

The **Multiprocessor Configuration Table** contains information needed when using RTEMS in a multiprocessor configuration. Many of the details associated with configuring a multiprocessor system are dependent on the multiprocessor communications layer provided by the user. The address of the **Multiprocessor Configuration Table** should be placed in the **Mp_tbl** entry in the primary **Configuration Table**. Further details regarding many of the entries in the **Multiprocessor Configuration Table** will be provided in the **Multiprocessing** chapter. The format of the **Multiprocessor Configuration Table** is defined in a C structure, and is given below:

```

struct mp_info {
    unsigned16  node;           /* local node number      */
    unsigned16  max_nodes;      /* number nodes in system */
    unsigned32  max_gobjects;   /* max global objects     */
    unsigned32  max_proxies;    /* max proxies            */
    mpci_table *Mpci_tbl;       /* MPCl table              */
};

```

node is a unique processor identifier and is used in routing messages between nodes in a multiprocessor configuration. Each processor must have a unique node number. RTEMS assumes that node numbers start at one and increase sequentially. This assumption can be used to advantage by the user-supplied MPCl layer. Typically, this requirement is made when the node numbers are used to calculate the address of inter-processor communication links. Zero should be avoided as a node number because some MPCl layers use node zero to represent broadcasted packets. Thus, it is recommended that node numbers start at one and increase sequentially.

max_nodes is the number of processor nodes in the system.

max_gobjects is the maximum number of global objects which can exist at any given moment in the entire system. If this parameter is not the same on all nodes in the system, then a fatal error

is generated to inform the user that the system is inconsistent.

max_proxies is the maximum number of proxies which can exist at any given moment on this particular node. A proxy is a substitute task control block which represent a task residing on a remote node when that task blocks on a remote object. Proxies are used in situations in which delayed interaction is required with a remote node.

Mpci_tbl is the address of the **Multiprocessor Communications Interface Table**. This table contains the entry points of user-provided functions which constitute the multiprocessor communications layer. This table must be provided in multiprocessor configurations with all entries configured. The format of this table and details regarding its entries can be found in the next section.

19.6 Multiprocessor Communications Interface Table

The format of this table is defined in a C structure, and is given below:

```
struct mpci_info {  
    proc_ptr init;          /* initialization procedure */  
    proc_ptr getpkt;        /* get packet procedure */  
    proc_ptr retpkt;        /* return packet procedure */  
    proc_ptr send;          /* packet send procedure */  
    proc_ptr receive;       /* packet receive procedure */  
};
```

init is the address of the entry point for the initialization procedure of the user supplied multiprocessor communications layer.

getpkt is the address of the entry point for the procedure called by **RTEMS** to obtain a packet from the user supplied multiprocessor communications layer.

retpkt is the address of the entry point for the procedure called by **RTEMS** to return a packet to the user supplied multiprocessor communications layer.

send is the address of the entry point for the procedure called by **RTEMS** to send an envelope to another node. This procedure is part of the user supplied multiprocessor communications layer.

receive is the address of the entry point for the procedure called by **RTEMS** to retrieve an envelope containing a message from another node. This procedure is part of the user supplied multiprocessor communications layer.

More information regarding the required functionality of these entry points is provided in the **Multiprocessor** chapter.

19.7 Determining Memory Requirements

Since memory is a critical resource in many real-time embedded systems, **RTEMS** was specifically designed to allow unused managers to be excluded from the run-time environment. This allows the application designer the flexibility to tailor **RTEMS** to most efficiently meet system requirements while still satisfying even the most stringent memory constraints. As result, the size of the **RTEMS** executive is application dependent. Appendix A provides a worksheet for calculating the memory requirements of a custom **RTEMS** run-time environment. To insure that enough memory is allocated for future versions of **RTEMS**, the application designer should round these memory requirements up. The following managers may be optionally excluded:

- *signal*
- *region*
- *dual ported memory*
- *I/O*
- *fatal error*
- *partition*
- *time*
- *semaphore*
- *message*
- *event*

RTEMS based applications must somehow provide memory for **RTEMS**' code and data space. Although **RTEMS**' data space must be in RAM, its code space can be located in either ROM or RAM. In addition, the user must allocate RAM for the **RTEMS RAM Workspace**. The size of this area is application dependent and can be calculated using the formula also provided Appendix A.


All **RTEMS** data variables and routine names used by **RTEMS** begin with the underscore (**_**) character. If **RTEMS** is linked with an application, then the application code should **NOT** contain any symbols which begin with the underscore character and an upper-case letter to avoid any naming conflicts. All **RTEMS** directive names should be treated as reserved words.

19.7.1 Sizing the RTEMS RAM Workspace

The **RTEMS RAM Workspace** is a user-specified block of memory reserved for use by **RTEMS**. The application should **NOT** modify this memory after it is cleared by the board support package. This area consists primarily of the **RTEMS** data structures whose exact size depends many values specified in the **Configuration Table**. In addition, task stacks and floating point context areas are dynamically allocated from the **RTEMS RAM Workspace**.

The starting address of the **RTEMS RAM Workspace** must be aligned on a four-byte boundary. Failure to properly align the workspace area will result in a fatal error condition.

A worksheet is provided in Appendix A to assist the user in calculating the minimum size of the **RTEMS RAM Workspace** for this application. The value calculated with this worksheet is the minimum value that should be specified as the **ram_size** parameter of the **Configuration Table**. The user is cautioned that future versions of **RTEMS** may not have the same memory requirements per object. Although the value calculated is sufficient for all processors supported by the current release of **RTEMS**, the user is advised to allocate somewhat more memory than the worksheet recommends to insure compatibility with future releases.



Multiprocessing Manager

20.1 Introduction

In multiprocessor real-time systems, new requirements, such as sharing data and global resources between processors, are introduced. This requires an efficient and reliable communications vehicle which allows all processors to communicate with each other as necessary. In addition, the ramifications of multiple processors affect each and every characteristic of a real-time system, almost always making them more complicated.

RTEMS addresses these issues by providing simple and flexible real-time multiprocessing capabilities. The executive easily lends itself to both tightly-coupled and loosely-coupled configurations of the target system hardware. In addition, **RTEMS** supports both homogeneous and heterogeneous target environments.

A major design goal of the **RTEMS** executive was to transcend the physical boundaries of the target hardware configuration. This goal is achieved by presenting to the application software a logical view of the target system where the boundaries between processor nodes are transparent. As a result, the application developer may designate objects such as tasks, queues, events, signals, semaphores, and memory blocks as global objects. These global objects may then be accessed by any task regardless of which processors the object and the accessing task may reside. **RTEMS** automatically determines that the object being accessed resides on another processor and performs the actions required to access the desired object. Simply stated, **RTEMS** allows the entire system, both hardware and software, to be viewed logically as a single system.

20.2 Background

RTEMS makes no assumptions regarding the connection media or the topology of a multiprocessor system. The tasks which compose a particular application can be spread among several processors. The application tasks can interact using a subset of the **RTEMS** directives as if they were on the same processor. These directives allow application tasks to exchange data, communicate, and synchronize regardless of which processor they reside upon.

The **RTEMS** multiprocessor execution model is multiple instruction streams with multiple data streams (MIMD). This execution model has each of the processors executing code independent of the other processors. Because of this parallelism, the application designer can more easily guarantee deterministic behavior.

By supporting heterogeneous environments, **RTEMS** allows the systems designer to select the most efficient processor for each subsystem of the application. Configuring **RTEMS** for a heterogeneous environment is no more difficult than for a homogeneous one. In keeping with **RTEMS** philosophy of providing transparent physical node boundaries, the minimal heterogeneous processing required is isolated in the **MPCI** layer.

20.2.1 Nodes

A processor in a **RTEMS** system is referred to as a node. Each node is assigned a unique non-zero node number by the application designer. **RTEMS** assumes that node numbers are assigned consecutively from one to **max_nodes**. The node number, **node**, and the maximum number of nodes, **max_nodes**, in a system are found in the **Multiprocessor Configuration Table**. The **max_nodes** field and the number of global objects, **num_gobjects**, is required to be the same on all nodes in a system.

The node number is used by **RTEMS** to identify each node when performing remote operations. Thus, the **Multiprocessor Communications Interface Layer (MPCI)** must be able to route messages based on the node number.

20.2.2 Global Objects

All **RTEMS** objects which are created with the **GLOBAL** option will be known on all other nodes. Global objects can be referenced from any node in the system, although certain directive specific restrictions (e.g. cannot delete a remote object) may apply. A task does not have to be global to perform operations involving remote objects. The distribution of tasks to processors is performed

during the application design phase. Dynamic task relocation is not supported by RTEMS.

20.2.3 Global Object Table

Every node in a multiprocessor system maintains two tables containing object information: a local object table and a global object table. The local object table on each node is unique and contains information for both local and global objects created on this node. The global object table contains information regarding all global objects in the system and therefore, is the same on every node.

Since each node must maintain an identical copy of the global object table, the maximum number of entries in each copy is determined by the **num_gobjects** parameter in the **Multiprocessor Configuration Table**. This parameter, as well as the **max_nodes** parameter, is required to be the same on all nodes. To maintain consistency among the table copies, every node in the system must be informed of the creation or deletion of a global object.

20.2.4 Remote Operations

When an application performs an operation on a remote global object, RTEMS must generate a Remote Request (RQ) message and send it to the appropriate node. After completing the requested operation, the remote node will build a Remote Response (RR) message and send it to the originating node. Messages generated as a side-effect of a directive (such as deleting a global task) are known as Remote Processes (RP) and do not require the receiving node to respond.

Other than taking slightly longer to execute directives on remote objects, the application is normally unaware of the location of the objects it acts upon. The exact amount of overhead required for a remote operation is dependent on the media connecting the nodes and, to a lesser degree, the efficiency of the user-provided MPCl routines.

The following shows the typical transaction sequence during a remote application:

- (1) The application issues a directive accessing a remote global object.
- (2) RTEMS determines the node on which the object resides.
- (3) RTEMS calls the user-provided MPCl routine GETPKT to obtain a packet in which to build a RQ message.

- (4) After building a message packet, RTEMS calls the user-provided MPCl routine **SEND** to transmit the packet to the node on which the object resides (referred to as the destination node).
- (5) The calling task is blocked until the RR message arrives, and control of the processor is transferred to another task.
- (6) The MPCl layer on the destination node senses the arrival of a packet (commonly in an ISR), and calls the **RTEMS mp_announce** directive. This directive readies the **Multiprocessing Server**.
- (7) The **Multiprocessing Server** calls the user-provided MPCl routine **RECEIVE**, performs the requested operation, builds an RR message, and returns it to the originating node.
- (8) The MPCl layer on the originating node senses the arrival of a packet (typically via an interrupt), and calls the **RTEMS mp_announce** directive. This directive readies the **Multiprocessing Server**.
- (9) The **Multiprocessing Server** calls the user-provided MPCl routine **RECEIVE**, readies the original requesting task, and blocks until another packet arrives. Control is transferred to the original task which then completes processing the directive.

If an uncorrectable error occurs in the user-provided MPCl layer, the fatal error handler should be invoked. **RTEMS** assumes the reliable transmission and reception of messages by the MPCl and makes no attempt to detect or correct errors.

20.2.5 Proxies

A **proxy** is an **RTEMS** data structure which resides on a remote node and is used to represent a task which must block as part of a remote operation. This action can occur as part of the **sm_p** and **q_receive** directives. If the object were local, the task's control block would be available for modification to indicate it was pending a message or semaphore. However, the task's control block resides only on the same node as the task. In this case, the remote node must allocate a proxy to represent the task until it can be readied.

The maximum number of proxies is defined in the **Multiprocessor Configuration Table**. Each node in a multiprocessor system may require a different number of proxies to be configured. The distribution of proxy control blocks is application dependent and is different from the distribution of tasks.

20.2.6 Multiprocessor Configuration Table

The **Multiprocessor Configuration Table** contains information needed by **RTEMS** when used in a multiprocessor system. This table is discussed in detail in a section in the previous chapter, **Multiprocessor Configuration Table**.

20.3 Multiprocessor Communications Interface Layer

The **Multiprocessor Communications Interface Layer (MPCI)** is a set of user-provided procedures which enable the nodes in a multiprocessor system to communicate with one another. These routines are invoked by **RTEMS** at various times in the generation and processing of remote requests. Interrupts are enabled when an MPCI procedure is invoked. It is assumed that if the execution mode and/or interrupt level are altered by the MPCI layer, that they will be restored prior to returning to **RTEMS**.

The MPCI layer is responsible for managing a pool of buffers called **packets** and for sending these packets between system nodes. Packet buffers contain the messages sent between the nodes. Typically, the MPCI layer will encapsulate the packet within an envelope which contains the information needed by the MPCI layer. The number of packets available is dependent on the MPCI layer implementation.

The entry points to the routines in the user's MPCI layer should be placed in the **Multiprocessor Communications Interface Table**. The user must provide entry points for each of the following table entries in a multiprocessor system:

init	initialize the MPCI
getpkt	obtain a packet buffer
retpkt	return a packet buffer
send	send a packet to another node
receive	called to get an arrived packet

A packet is sent by **RTEMS** in each of the following situations:

- *an RQ is generated on an originating node;*
- *an RR is generated on a destination node;*
- *a global object is created;*
- *a global object is deleted;*
- *a local task blocked on a remote object is deleted;*
- *during system initialization to check for system consistency.*

The arrival of a packet at a node may generate an interrupt. If it does not, the real-time clock ISR can check for the arrival of a packet. In any case, the **mp_announce** directive must be called to announce the arrival of a packet. After exiting the ISR, control will be passed to the Multiprocessing Server to process the packet. The Multiprocessing Server will call the **getpkt** entry to obtain a packet buffer and the **receive** entry to copy the message into the buffer obtained.

20.3.1 INIT

The INIT component of the user-provided MPCl layer is called as part of the **init_exec** directive to initialize the MPCl layer and associated hardware. It is invoked immediately after all of the device drivers have been initialized. This component should be prototyped as follows:

```
void init( conf_tbl )
config_table *conf_tbl;
```

where **conf_tbl** is the address of the user's Configuration Table. Operations on global objects cannot be performed until this component is invoked. The INIT component is invoked only once in the life of any system. If the MPCl layer cannot be successfully initialized, the fatal error manager should be invoked.

One of the primary functions of the MPCl layer is to provide the executive with packet buffers. The INIT routine must create and initialize a pool of packet buffers. There must be enough packet buffers so RTEMS can obtain one whenever needed.

20.3.2 GETPKT

The GETPKT component of the user-provided MPCl layer is called when RTEMS must obtain a packet buffer to send or broadcast a message. This component should be prototyped as follows:

```
void getpkt( pkt )
unsigned8 **pkt;
```

where **pkt** is the address of a pointer to a packet. This routine always succeeds and, upon return, **pkt** will contain the address of a packet. If for any reason, a packet cannot be successfully obtained, then the fatal error manager should be invoked.

RTEMS has been optimized to avoid the need for obtaining a packet each time a message is sent or broadcast. For example, **RTEMS** sends response messages (RR) back to the originator in the same packet in which the request message (RQ) arrived.

20.3.3 RETPKT

The **RETPKT** component of the user-provided MPCPI layer is called when **RTEMS** needs to release a packet to the free packet buffer pool. This component should be prototyped as follows:

```
void retpkt( pkt )
unsigned8 *pkt;
```

where **pkt** is the address of a packet. If the packet cannot be successfully returned, the fatal error manager should be invoked.

20.3.4 RECEIVE

The **RECEIVE** component of the user-provided MPCPI layer is called when **RTEMS** needs to obtain a packet which has previously arrived. This component should be prototyped as follows:

```
void receive( pkt )
unsigned8 **pkt;
```

where **pkt** is a pointer to the address of a packet to place the message from another node. If a message is available, then **pkt** will contain the address of the message from another node. If no messages are available, this entry **pkt** should contain **NULL_PACKET**.

20.3.5 SEND

The **SEND** component of the user-provided MPCPI layer is called when **RTEMS** needs to send a packet containing a message to another node. This component should be prototyped as follows:

```
void send( node, pkt, pkt_length )
unsigned16 node;
unsigned8 *pkt;
unsigned32 pkt_length;
```


where **node** is the node number of the destination, **pkt** is the address of a packet which containing a message, and **pkt_length** is the length of the message in bytes. If the packet cannot be successfully sent, the fatal error manager should be invoked.

If **node** is set to zero, the packet is to be broadcasted to all other nodes in the system. Although some MPCPI layers will be built upon hardware which support a broadcast mechanism, others may be required to generate a copy of the packet for each node in the system.


Many MPCPI layers use the **pkt_length** to avoid sending unnecessary data. This is especially useful if the media connecting the nodes is relatively slow.

20.3.6 Supporting Heterogeneous Environments

Developing an MPCPI layer for a heterogeneous system requires a thorough understanding of the differences between the processors which comprise the system. One difficult problem is the varying data representation schemes used by different processor types. The most pervasive data representation problem is the order of the bytes which compose a data entity. Processors which place the least significant byte at the smallest address are classified as **little endian** processors. Little endian byte-ordering is shown below:



Conversely, processors which place the most significant byte at the smallest address are classified as **big endian** processors. Big endian byte-ordering is shown below:



Unfortunately, sharing a data structure between big endian and little endian processors requires translation into a common endian format. An application designer typically chooses the common endian format to minimize conversion overhead.

Another issue in the design of shared data structures is the alignment of data structure elements. Alignment is both processor and compiler implementation dependent. For example, some processors allow data elements to begin on any address boundary, while others impose restrictions. Common restrictions are that data elements must begin on either an even address or on a long word

boundary. Violation of these restrictions may cause an exception or impose a performance penalty.

Other issues which commonly impact the design of shared data structures include the representation of floating point numbers, bit fields, decimal data, and character strings. In addition, the representation method for negative integers could be one's or two's complement. These factors combine to increase the complexity of designing and manipulating data structures shared between processors.

RTEMS addressed these issues in the design of the packets used to communicate between nodes. The **RTEMS** packet format is designed to allow the **MPCI** layer to perform all necessary conversion without burdening the developer with the details of the **RTEMS** packet format. As a result, the **MPCI** layer must be aware of the following:

- *All packets must begin on a long-word boundary.*
- *Packets are composed of both **RTEMS** and application data. All **RTEMS** data is unsigned32 and is located in the first **MIN_HETERO_CONV** unsigned32's of the packet.*
- *The **RTEMS** data component of the packet must be in native endian format. Endian conversion may be performed by either the sending or receiving **MPCI** layer.*
- ***RTEMS** makes no assumptions regarding the application data component of the packet.*

20.4 Operations

20.4.1 Announcing a Packet

The **mp_announce** directive is called by the **MPCI** layer to inform **RTEMS** that a packet has arrived from another node. This directive can be called from an interrupt service routine or from within a polling routine.

20.5 Directives

This section details the additional directives required to support **RTEMS** in a multiprocessor configuration. A subsection is dedicated to each of this manager's directives and describes the calling sequence, related constants, usage, and status codes.

20.5.1 MP_ANNOUNCE - Announce the arrival of a packet

CALLING SEQUENCE:

`void mp_announce()`

INPUT: NONE

OUTPUT: NONE

DIRECTIVE STATUS CODES: NONE

DESCRIPTION:

This directive informs **RTEMS** that a multiprocessing communications packet has arrived from another node. This directive is called by the user-provided **MPCI**, and is only used in multiprocessor configurations.

NOTES:

This directive is typically called from an ISR.

This directive will not cause the calling task to be preempted.

This directive does not generate activity on remote nodes.

A

Memory Requirements

A.1 Data Space Requirements

RTEMS requires a small amount of memory for its private variables. This data area must be in RAM and is separate from the **RTEMS RAM Workspace**. The following table illustrates the data space required for all configurations of **RTEMS**:

Component	Size
Data Space	800

A.2 Minimum and Maximum Code Space Requirements

A maximum configuration of **RTEMS** includes the core and all managers including the multiprocessing manager. Conversely, a minimum configuration of **RTEMS** includes only the core and the following managers: initialization, task, interrupt and fatal error. The following table illustrates the code space required by these configurations of **RTEMS**:

Configuration	Size
Minimum	7,296
Maximum	19,868

A.3 Code Space Worksheet

Component	Included	Not Included	Size
Core	2864	n/a	
Initialization	1184	n/a	
Task	3080	n/a	
Interrupt	52	n/a	
Fatal Error	32	n/a	
Time	2424	8	
Semaphore	1064	8	
Message	1920	8	
Event	604	8	
Signal	452	8	
Partition	1124	8	
Region	1072	8	
Dual Ported Memory	720	8	
I/O	112	8	
Multiprocessing	3164	12	
Total Code Space Requirements			

A.4 RTEMS RAM Workspace Worksheet

Description	Equation	Bytes Required
max_tasks	* 268 =	
max_semaphores	* 56 =	
max_timers	* 36 =	
max_queues	* 72 =	
max_messages	* 24 =	
max_regions	* 96 =	
max_partitions	* 52 =	
max_dpmem	* 32 =	
FP Tasks	* 108 =	
Task Stacks	=	

Total Single Processor Requirements	
--	--

max_nodes	* 48 =	
max_objects	* 12 =	
max_proxies	* 76 =	

Total Multiprocessing Requirements	
---	--

Fixed System Requirements	4428
Total Single Processor Requirements	
Total Multiprocessor Requirements	
Minimum Bytes for RTEMS Workspace	

B

DIRECTIVE STATUS CODES

CONSTANT	CODE	DESCRIPTION
SUCCESSFUL	0	successful completion
E_EXITTED	1	returned from a task
E_NOMP	2	multiprocessing not configured
E_NAME	3	invalid object name
E_ID	4	invalid object id
E_TOOMANY	5	too many
E_TIMEOUT	6	timed out waiting
E_DELETE	7	object was deleted while waiting
E_SIZE	8	invalid specified size
E_ADDRESS	9	invalid address specified
E_NUMBER	10	number was invalid
E_NOTDEFINED	11	item not initialized
E_INUSE	12	resources outstanding
E_UNSATISFIED	13	request not satisfied
E_STATE	14	task is in wrong state
E_ALREADY	15	task already in state
E_SELF	16	illegal for calling task
E_REMOTE	17	illegal for remote object
E_CALLED	18	invalid environment
E_PRIORITY	19	invalid task priority
E_CLOCK	20	invalid time buffer
E_NODE	21	invalid node id
E_NOTCONFIGURED	22	directive not configured
E_NOTIMPLEMENTED	23	directive not implemented

C

HEADER FILES

C.1 Header File Usage

RTEMS is distributed with a single C header file, **rtems.h**, which contains all structure and macro definitions needed to utilize **RTEMS**. This header file must be included by all C modules which utilize **RTEMS**. It contains type definitions, NULL pointer definitions, directive status codes, **RTEMS** data structures, flag and mode definitions, event and signal sets, task manager related definitions, and the prototypes for the entries in the C language interface library.

C.2 **rtems.h**

```
/* rtems.h
 *
 * This include file contains information about the
 * executive that is needed by the application.
 */

#ifndef __RTEMS_h
#define __RTEMS_h

/* processor dependent type definitions */

typedef unsigned char  unsigned8; /* unsigned 8-bit value */
typedef unsigned short unsigned16; /* unsigned 16-bit value */
typedef unsigned int   unsigned32; /* unsigned 32-bit value */

/* RTEMS dependent type definitions */

typedef void          task;          /* RTEMS task          */
typedef task          (*task_ptr)(); /* task pointer        */
typedef void          (*proc_ptr)(); /* procedure ptr       */
typedef void          (*asr_ptr)();  /* ASR pointer         */
typedef unsigned32    obj_name;      /* object name         */
typedef unsigned32    obj_id;        /* object id           */
typedef unsigned32    dir_status;    /* directive status    */
typedef unsigned32    task_mode;     /* task mode           */
typedef unsigned32    task_pri;      /* task priority       */
```

```

typedef unsigned32 interval;          /* tick interval */
typedef unsigned32 event_set;         /* event set */
typedef unsigned32 signal_set;        /* signal set */

typedef struct t_ctlblk t_cb;         /* task control */
typedef struct time_info time_buffer; /* date/time info */
typedef struct ext_info ext_table;    /* user extensions */
typedef struct driver_info driver_table; /* driver table */
typedef struct itasks_info itasks_table; /* init tasks */
typedef struct mpci_info mpci_table;  /* MPCI entries */
typedef struct config_info config_table; /* config table */
typedef struct mp_info mp_table;      /* mp config table */

/* RTEMS directive completion statuses */

#define SUCCESSFUL 0 /* successful completion */
#define E_EXITED 1 /* returned from a task */
#define E_NOMP 2 /* multiprocessing not configured */
#define E_NAME 3 /* invalid object name */
#define E_ID 4 /* invalid object id */
#define E_TOOMANY 5 /* too many */
#define E_TIMEOUT 6 /* timed out waiting */
#define E_DELETE 7 /* object was deleted while waiting */
#define E_SIZE 8 /* invalid specified size */
#define E_ADDRESS 9 /* invalid address specified */
#define E_NUMBER 10 /* number was invalid */
#define E_NOTDEFINED 11 /* item not initialized */
#define E_INUSE 12 /* resources outstanding */
#define E_UNSATISFIED 13 /* request not satisfied */
#define E_STATE 14 /* task is in wrong state */
#define E_ALREADY 15 /* task already in state */
#define E_SELF 16 /* illegal for calling task */
#define E_REMOTE 17 /* illegal for remote object */
#define E_CALLED 18 /* invalid environment */
#define E_PRIORITY 19 /* invalid task priority */
#define E_CLOCK 20 /* invalid date/time */
#define E_NODE 21 /* invalid node id */
#define E_NOTCONFIGURED 22 /* directive not configured */
#define E_NOTIMPLEMENTED 23 /* directive not implemented */

/* Task states */

#define TS_READY 0x000 /* ready to run */
#define TS_DORMANT 0x001 /* created, not started */
#define TS_SUSPEND 0x002 /* wait to be resumed */
#define TS_TRANSIENT 0x004 /* task in transition */
#define TS_DELAY 0x008 /* wait for timeout */
#define TS_WAITSEGMENT 0x010 /* wait for segment */
#define TS_WAITMESSAGE 0x020 /* wait for message */
#define TS_WAITEVENT 0x040 /* wait for event */
#define TS_WAITSEMAPHORE 0x080 /* wait for semaphore */
#define TS_WAITTIME 0x100 /* wait for date and time */
#define TS_WAITRPCREPLY 0x200 /* wait for rpc reply */

/* defaults for attributes or options */

#define DEFAULTS 0x00000000 /* all default options */

```

```

/* attribute constants */

#define NOFP      0x00000000    /* no floating point unit */
#define FP        0x00000001    /* floating point unit */

#define LOCAL     0x00000000    /* local object */
#define GLOBAL    0x00000002    /* global object */

#define FIFO      0x00000000    /* process in FIFO order */
#define PRIORITY  0x00000004    /* process by priority */

#define NOLIMIT   0x00000000    /* do not place a limit */
#define LIMIT     0x00000008    /* limit queue entries */

/* options constants */

#define WAIT      0x00000000    /* wait for completion */
#define NOWAIT    0x00000001    /* do not wait on resource */

#define EV_ALL    0x00000000    /* wait for all event(s) */
#define EV_ANY    0x00000002    /* wait for any event(s) */

/* mask constants */

#define PREEMPTMODE 0x00000100 /* preemption bit */
#define TSLICEMODE  0x00000200 /* tslice bit */
#define ASRMODE     0x00000400 /* ASR enable bit */
#define INTRMODE    0x00000007 /* interrupt mode bits */

/* mode constants */

#define PREEMPT     0x00000000 /* enable preemption */
#define NOPREEMPT   0x00000100 /* disable preemption */

#define NOTSLICE    0x00000000 /* disable timeslicing */
#define TSLICE      0x00000200 /* enable timeslicing */

#define ASR         0x00000000 /* enable ASR */
#define NOASR       0x00000400 /* disable ASR */

#define INTR(level) (level&0x7) /* interrupt level */

/* null constants */

#define NULL_PACKET      (unsigned8 *) 0
#define NULL_DP_ADDRESS  (unsigned8 *) 0
#define NULL_TASK        (task_ptr) 0
#define NULL_ASR         (asr_ptr) 0
#define NULL_DRIVER      (proc_ptr) 0
#define NULL_DRIVER_TABLE (driver_table *) 0
#define NULL_ITASKS_TABLE (itasks_table *) 0
#define NULL_EXT_TABLE    (ext_table *) 0
#define NULL_EXTENSION    (proc_ptr) 0
#define NULL_MPCI_TABLE   (mpci_table *) 0
#define NULL_MP_TABLE     (mp_table *) 0

/* Identification constants */
#define ALL_NODES 0 /* search all nodes */
#define OTHER_NODES 0xffffffff /* search all except local */
#define LOCAL_NODE 0xffffffff /* only local node */
#define WHO_AM_I 0 /* calling task */

```

/* miscellaneous constants */

```
#define CURRENT      0      /* current mode/priority */
#define NOTIMEOUT    0      /* wait indefinitely */
#define SELF         0      /* current task */
#define YIELD        0      /* yield CPU (tm_wkafter) */
#define MIN_PRIORITY 1      /* highest task priority */
#define MAX_PRIORITY 255    /* lowest task priority */
#define MIN_STK_SIZE 256    /* minimum task stack size */
```

/* definitions for event set */

```
#define EVENT_0      0x00000001
#define EVENT_1      0x00000002
#define EVENT_2      0x00000004
#define EVENT_3      0x00000008
#define EVENT_4      0x00000010
#define EVENT_5      0x00000020
#define EVENT_6      0x00000040
#define EVENT_7      0x00000080
#define EVENT_8      0x00000100
#define EVENT_9      0x00000200
#define EVENT_10     0x00000400
#define EVENT_11     0x00000800
#define EVENT_12     0x00001000
#define EVENT_13     0x00002000
#define EVENT_14     0x00004000
#define EVENT_15     0x00008000
#define EVENT_16     0x00010000
#define EVENT_17     0x00020000
#define EVENT_18     0x00040000
#define EVENT_19     0x00080000
#define EVENT_20     0x00100000
#define EVENT_21     0x00200000
#define EVENT_22     0x00400000
#define EVENT_23     0x00800000
#define EVENT_24     0x01000000
#define EVENT_25     0x02000000
#define EVENT_26     0x04000000
#define EVENT_27     0x08000000
#define EVENT_28     0x10000000
#define EVENT_29     0x20000000
#define EVENT_30     0x40000000
#define EVENT_31     0x80000000
```

/* definitions for signal set */

```
#define SIGNAL_0      0x00000001
#define SIGNAL_1      0x00000002
#define SIGNAL_2      0x00000004
#define SIGNAL_3      0x00000008
#define SIGNAL_4      0x00000010
#define SIGNAL_5      0x00000020
#define SIGNAL_6      0x00000040
#define SIGNAL_7      0x00000080
#define SIGNAL_8      0x00000100
#define SIGNAL_9      0x00000200
#define SIGNAL_10     0x00000400
#define SIGNAL_11     0x00000800
#define SIGNAL_12     0x00001000
```

```

#define SIGNAL_13    0x00002000
#define SIGNAL_14    0x00004000
#define SIGNAL_15    0x00008000
#define SIGNAL_16    0x00010000
#define SIGNAL_17    0x00020000
#define SIGNAL_18    0x00040000
#define SIGNAL_19    0x00080000
#define SIGNAL_20    0x00100000
#define SIGNAL_21    0x00200000
#define SIGNAL_22    0x00400000
#define SIGNAL_23    0x00800000
#define SIGNAL_24    0x01000000
#define SIGNAL_25    0x02000000
#define SIGNAL_26    0x04000000
#define SIGNAL_27    0x08000000
#define SIGNAL_28    0x10000000
#define SIGNAL_29    0x20000000
#define SIGNAL_30    0x40000000
#define SIGNAL_31    0x80000000

/* notepad location definitions */

#define NOTEPAD_0      0
#define NOTEPAD_1      1
#define NOTEPAD_2      2
#define NOTEPAD_3      3
#define NOTEPAD_4      4
#define NOTEPAD_5      5
#define NOTEPAD_6      6
#define NOTEPAD_7      7
#define NOTEPAD_8      8
#define NOTEPAD_9      9
#define NOTEPAD_10     10
#define NOTEPAD_11     11
#define NOTEPAD_12     12
#define NOTEPAD_13     13
#define NOTEPAD_14     14
#define NOTEPAD_15     15

/* Multiprocessing constants */

#define MIN_PKT_SIZE 64      /* MPCPI layer must support */
                             /* packets = this size */

/* The following constant defines the number of unsigned32's
 * in a pkt which must be converted to native format in a
 * heterogeneous system. In packets longer than
 * MIN_HETERO_CONV unsigned32's, the "extra" data is a
 * message buffer.
 */

#define MIN_HETERO_CONV 7

/* macros to access sub-fields in an object id */

#define _Node( id )      (id > 16)      /* MSW = local node id */
#define _Object( id )    (id & 0xffff) /* LSW = object's id */

#define Build_name( N, C1, C2, C3, C4 ) \
    *(N) = Ret_name( C1,C2,C3,C4 )

```

```

#define Ret_name( c1, c2, c3, c4 ) ( c1 < 24 | c2 < 16 | c3 < 8 |
c4 )

/* structures */

/* date/time */

struct time_info {
    unsigned16 year;          /* year, A.D. */
    unsigned8 month;          /* month, 1 - 12 */
    unsigned8 day;            /* day, 1 - 31 */
    unsigned16 hour;          /* hour, 0 - 23 */
    unsigned8 minute;         /* minute, 0 - 59 */
    unsigned8 second;         /* second, 0 - 59 */
    unsigned32 ticks;         /* elapsed ticks between secs */
};

/* I80386 registers */

struct registers {
    unsigned32 _eip;          /* task's stack pointer; other */
                                /* registers pushed on stack */
                                /* using "pushad" */
};

/* task control block */

struct t_ctlblk {
    t_cb *next;               /* pointer to next TCB */
    t_cb *prev;               /* pointer to previous TCB */
    obj_id tid;               /* task identification */
    obj_name name;            /* task name */
    unsigned32 state;         /* task state */
    task_pri priority;        /* current task priority */
    unsigned8 reserved[128];  /* reserved for RTEMS */
    task_mode mode;           /* current task mode */
    unsigned32 attributes;    /* task attributes */
    struct registers Regs;    /* I80386 register save area */
    unsigned8 *fp_context;    /* pointer to FP context area */
    unsigned32 notepad[16];   /* executive notepads */
    unsigned8 *extension;     /* pointer TCB extension */
};

/* init task(s) information */

struct itasks_info {
    obj_name name;            /* task name */
    unsigned32 stksize;       /* task stack size */
    task_pri priority;        /* task priority */
    unsigned32 attributes;    /* task attributes */
    task_ptr entry_point;     /* task entry point */
    task_mode mode;           /* task initial mode */
    unsigned32 arg;           /* task argument */
};

/* device driver table */

struct driver_info {
    proc_ptr init;            /* initialization procedure */
    proc_ptr open;            /* open request procedure */
    proc_ptr close;           /* close request procedure */
};

```



```

    proc_ptr    read;          /* read request procedure */
    proc_ptr    write;         /* write request procedure */
    proc_ptr    cntrl;         /* special requests procedure */
    unsigned32  reserved1;     /* reserved for RTEMS */
    unsigned32  reserved2;     /* reserved for RTEMS */
};

/* user extensions table */

struct ext_info {
    proc_ptr    tcreate;       /* tcreate user extension */
    proc_ptr    tstart;        /* tstart user extension */
    proc_ptr    trestart;      /* trestart user extension */
    proc_ptr    tdelete;       /* tdelete user extension */
    proc_ptr    tswitch;       /* tswitch user extension */
    proc_ptr    taskexitted;   /* task exited user extension */
    proc_ptr    fatal;         /* fatal error user extension */
};

/* multiprocessor communications interface table */

struct mpci_info {
    proc_ptr    init;          /* initialization procedure */
    proc_ptr    getpkt;        /* get packet procedure */
    proc_ptr    retpkt;        /* return packet procedure */
    proc_ptr    send;          /* packet send procedure */
    proc_ptr    receive;       /* packet receive procedure */
};

/* multiprocessor configuration table */

struct mp_info {
    unsigned16  node;          /* local node number */
    unsigned16  max_nodes;     /* number nodes in system */
    unsigned32  max_gobjects;  /* maximum global objects */
    unsigned32  max_proxies;   /* maximum proxies */
    mpci_table *mpci_tbl;      /* MPCl table */
};

/* configuration table */

struct config_info {
    unsigned32  exec_ram;       /* RTEMS work area */
    unsigned32  ram_size;       /* RTEMS work area size */
    unsigned16  max_tasks;      /* max number tasks */
    unsigned16  max_semaphores; /* max number semaphores */
    unsigned16  max_timers;     /* max number timers */
    unsigned16  max_queues;     /* max number queues */
    unsigned16  max_messages;   /* max number messages */
    unsigned16  max_regions;    /* max number regions */
    unsigned16  max_partitions; /* max number partitions */
    unsigned16  max_dpmems;     /* max dp memory areas */
    unsigned16  ms_tick;        /* number ms in a tick */
    unsigned16  tslice;         /* ticks in a timeslice */
    unsigned32  num_itasks;     /* number init tasks */
    itasks_table *itasks_tbl;   /* init tasks table */
    unsigned32  num_devices;    /* number device drivers */
    driver_table *drv_tbl;      /* device driver table */
    ext_table   *ext_tbl;       /* user extension table */
    mp_table    *mp_tbl;        /* MP config table */
};

```

```
/* C Interface Definitions */
```

```
/*  
 * The following macros are used to map user directive  
 * names to the C Interface entry points. The exact  
 * names of the C Interface entry points may vary  
 * between compilers.  
 */
```

```
#define      as_catch      C_As_catch  
#define      as_send      C_As_send  
  
#define      de_close     C_De_close  
#define      de_cntrl     C_De_cntrl  
#define      de_init      C_De_init  
#define      de_open      C_De_open  
#define      de_read      C_De_read  
#define      de_write     C_De_write  
  
#define      dp_2external C_Dp_2external  
#define      dp_2internal C_Dp_2internal  
#define      dp_create    C_Dp_create  
#define      dp_delete    C_Dp_delete  
#define      dp_ident     C_Dp_ident  
  
#define      ev_receive   C_Ev_receive  
#define      ev_send      C_Ev_send  
  
#define      init_exec    C_Init_exec  
  
#define      pt_create    C_Pt_create  
#define      pt_delete    C_Pt_delete  
#define      pt_getbuf    C_Pt_getbuf  
#define      pt_ident     C_Pt_ident  
#define      pt_retbuf    C_Pt_retbuf  
  
#define      q_broadcast  C_Q_broadcast  
#define      q_create     C_Q_create  
#define      q_delete     C_Q_delete  
#define      q_flush      C_Q_flush  
#define      q_ident      C_Q_ident  
#define      q_receive    C_Q_receive  
#define      q_send       C_Q_send  
#define      q_urgent     C_Q_urgent  
  
#define      rn_create    C_Rn_create  
#define      rn_delete    C_Rn_delete  
#define      rn_getseg    C_Rn_getseg  
#define      rn_ident     C_Rn_ident  
#define      rn_retseg    C_Rn_retseg  
  
#define      sm_create    C_Sm_create  
#define      sm_delete    C_Sm_delete  
#define      sm_ident     C_Sm_ident  
#define      sm_p         C_Sm_p  
#define      sm_v         C_Sm_v  
  
#define      t_create     C_T_create  
#define      t_delete     C_T_delete  
#define      t_getnote    C_T_getnote
```

```

#define      t_ident      C_T_ident
#define      t_mode       C_T_mode
#define      t_restart    C_T_restart
#define      t_resume     C_T_resume
#define      t_setpri     C_T_setpri
#define      t_setnote    C_T_setnote
#define      t_start      C_T_start
#define      t_suspend    C_T_suspend

#define      tm_cancel    C_Tm_cancel
#define      tm_evafter   C_Tm_evafter
#define      tm_evevery   C_Tm_evevery
#define      tm_evwhen    C_Tm_evwhen
#define      tm_get       C_Tm_get
#define      tm_set       C_Tm_set
#define      tm_tick      C_Tm_tick
#define      tm_wkafter   C_Tm_wkafter
#define      tm_wkwhen    C_Tm_wkwhen

#define      k_fatal      C_K_fatal

#define      mp_announce  C_Mp_announce

/* C Language Directive Prototypes */

dir_status C_As_catch();
dir_status C_As_send();

dir_status C_De_close();
dir_status C_De_cntrl();
dir_status C_De_init();
dir_status C_De_open();
dir_status C_De_read();
dir_status C_De_write();

dir_status C_Dp_2external();
dir_status C_Dp_2internal();
dir_status C_Dp_create();
dir_status C_Dp_delete();
dir_status C_Dp_ident();

dir_status C_Ev_receive();
dir_status C_Ev_send();

void C_Init_exec();

dir_status C_Pt_create();
dir_status C_Pt_delete();
dir_status C_Pt_getbuf();
dir_status C_Pt_ident();
dir_status C_Pt_retbuf();

dir_status C_Q_broadcast();
dir_status C_Q_create();
dir_status C_Q_delete();
dir_status C_Q_flush();
dir_status C_Q_ident();
dir_status C_Q_receive();
dir_status C_Q_send();
dir_status C_Q_urgent();

```

```

dir_status C_Rn_create();
dir_status C_Rn_delete();
dir_status C_Rn_getseg();
dir_status C_Rn_ident();
dir_status C_Rn_retseg();

dir_status C_Sm_create();
dir_status C_Sm_delete();
dir_status C_Sm_ident();
dir_status C_Sm_p();
dir_status C_Sm_v();

dir_status C_T_create();
dir_status C_T_delete();
dir_status C_T_getnote();
dir_status C_T_ident();
dir_status C_T_mode();
dir_status C_T_restart();
dir_status C_T_resume();
dir_status C_T_setpri();
dir_status C_T_setnote();
dir_status C_T_start();
dir_status C_T_suspend();

dir_status C_Tm_cancel();
dir_status C_Tm_evafter();
dir_status C_Tm_evevery();
dir_status C_Tm_evwhen();
dir_status C_Tm_get();
dir_status C_Tm_set();
dir_status C_Tm_tick();
dir_status C_Tm_wkafter();
dir_status C_Tm_wkwhen();

void C_K_fatal();

void C_Mp_announce();

#endif
/* end of include file */

```

EXAMPLE APPLICATION

```

/* example.c
 *
 * This file contains an example of a simple RTEMS
 * application. It contains a Configuration Table, a
 * user initialization task, and a simple task.
 *
 * This example assumes that a board support package exists
 * and invokes the in_executive() directive. The board
 * support package also provides the function Task_exitted()
 * for the task exited user-supplied routine.
 */

#include "rtems.h"

task init_task();

struct itasks_info init_task = {
    { 0x61626300, /* init task name "ABC" */
      1024, /* init task stack size */
      1, /* init task priority */
      DEFAULTS, /* init task attributes */
      init_task, /* init task entry point */
      TSLICE, /* init task initial mode */
      0, /* init task argument */
    }
};

struct config_tbl config_tbl = {
    0x0f0000, /* exective RAM work area */
    65536, /* exective RAM size */
    2, /* maximum tasks */
    0, /* maximum semaphores */
    0, /* maximum timers */
    0, /* maximum message queues */
    0, /* maximum messages */
    0, /* maximum regions */
    0, /* maximum partitions */
    0, /* maximum dp memory areas */
    10, /* number of ms in a tick */
    1, /* num of ticks in a timeslice */
};

```

```

1,          /* number of user init tasks */
init_task_tbl, /* user init task(s) table */
0,          /* number of device drivers */
NULL_DRIVER_TABLE, /* ptr to driver address table */
NULL_EXT_TABLE, /* ptr to extension table */
NULL_MP_TABLE, /* ptr to MP config table */
};

task user_application();

#define USER_APP_NAME 1 /* any 32-bit name; unique helps */

task init_task()
{
    obj_id tid;

    /* example assumes SUCCESSFUL return value */
    t_create( USER_APP_NAME, 1, 1024, NOPREEMPT, FP, &tid );
    t_start( app_tid, user_application, 0 );
    t_delete( SELF );
}

task user_application()
{
    /* application specific initialization goes here */

    while ( 1 ) { /* infinite loop */
        /* APPLICATION CODE GOES HERE
        *
        * This code will typically include at least one
        * directive which causes the calling task to
        * give up the processor.
        */
    }
}

```

E

GLOSSARY

active	A term used to describe an object which has been created by an application.
application	In this document, software which makes use of RTEMS.
ASR	see Asynchronous Signal Routine.
asynchronous	Not related in order or timing to other occurrences in the system.
Asynchronous Signal Routine	Similar to a hardware interrupt except that it is associated with a task and is run in the context of a task. The directives provided by the signal manager are used to service signals.
awakened	A term used to describe a task that has been unblocked and may be scheduled to the CPU.
big endian	A data representation scheme in which the bytes composing a numeric value are arranged such that the most significant byte is at the lowest address.
bit-mapped	A data encoding scheme in which each bit in a variable is used to represent something different. This makes for compact data representation.
block	A physically contiguous area of memory.
blocked	The task state entered by a task which has been previously started and cannot continue execution until the reason for waiting has been satisfied.

broadcast	To simultaneously send a message to a logical set of destinations.
BSP	see Board Support Package.
Board Support Package	A collection of device initialization and control routines specific to a particular type of board or collection of boards.
buffer	A fixed length block of memory allocated from a partition.
Central Processing Unit	This term is equivalent to the terms processor and microprocessor.
chain	A data structure which allows for efficient dynamic addition and removal of elements. It differs from an array in that it is not limited to a predefined size.
coalesce	The process of merging adjacent holes into a single larger hole. Sometimes this process is referred to as garbage collection.
Configuration Table	A table which contains information used to tailor RTEMS for a particular application.
context	All of the processor registers and operating system data structures associated with a task.
context switch	Alternate term for task switch. Taking control of the processor from one task and transferring it to another task.
control block	A data structure used by the executive to define and control an object.
core	When used in this manual, this term refers to the internal executive utility functions. In the interest of application portability, the core of the executive should not be used directly by applications.
CPU	An acronym for Central Processing Unit.
critical section	A section of code which must be executed indivisibly.
CRT	An acronym for Cathode Ray Tube. Normally used in reference to the man-machine interface.

device	A peripheral used by the application that requires special operation software. See also device driver.
device driver	Control software for special peripheral devices used by the application.
directives	RTEMS' provided routines that provide support mechanisms for real-time applications.
dispatch	The act of loading a task's context onto the CPU and transferring control of the CPU to that task.
dormant	The state entered by a task after it is created and before it has been started.
Driver Address Table	A table which contains the entry points for each of the configured device drivers.
dual-ported	A term used to describe memory which can be accessed at two different addresses.
embedded	An application that is delivered as a hidden part of a larger system. For example, the software in a fuel-injection control system is an embedded application found in many late-model automobiles.
envelope	A buffer provided by the MPCI layer to RTEMS which is used to pass messages between nodes in a multiprocessor system. It typically contains routing information needed by the MPCI. The contents of an envelope are referred to as a packet.
entry point	The address at which a function or task begins to execute. In C, the symbolic entry point of a function is the function's name.
events	A method for task communication and synchronization. The directives provided by the event manager are used to service events.
exception	A synonym for interrupt.
executing	The task state entered by a task after it has been given control of the CPU.

executive	In this document, this term is used to referred to RTEMS. Commonly, an executive is a small real-time operating system used in embedded systems.
exported	An object known by all nodes in a multiprocessor system. An object created with the GLOBAL attribute set will be exported.
external address	The address used to access dual-ported memory by all the nodes in a system which do not own the memory.
FIFO	An acronym for First In First Out.
First In First Out	A discipline for manipulating entries in a data structure.
floating point coprocessor	A component used in computer systems to enhance performance in mathematically intensive situations. It is typically viewed as a logical extension of the primary processor.
freed	A resource that has been released by the application to RTEMS.
global	An object that has been created with the GLOBAL attribute set and exported to all nodes in a multiprocessor system.
handler	The equivalent of a manager, except that it is internal to RTEMS and forms part of the core. A handler is a collection of routines which provide a related set of functions. For example, there is a handler used by RTEMS to manage all objects.
heap	A data structure used to dynamically allocate and deallocate variable sized blocks of memory.
heterogeneous	A multiprocessor computer system composed of dissimilar processors.
homogeneous	A multiprocessor computer system composed of one type of processor.
ID	An RTEMS assigned identification tag used to access an active object.

IDLE task	A special low priority task which assumes control of the CPU when no other task is able to execute.
Instruction Pointer	A hardware register containing the address of the current instruction.
interface	A specification of the methodology used to connect multiple independent subsystems.
internal address	The address used to access dual-ported memory by the node which owns the memory.
interrupt	A hardware facility that causes the CPU to suspend execution, save its status, and transfer control to a specific location.
interrupt level	A mask used to by the CPU to determine which pending interrupts should be serviced. If a pending interrupt is below the current interrupt level, then the CPU does not recognize that interrupt.
Interrupt Service Routine	An ISR is invoked by the CPU to process a pending interrupt.
I/O	An acronym for Input/Output.
IP	An acronym for Instruction Pointer.
ISR	An acronym for Interrupt Service Routine.
kernel	In this document, this term is used as a synonym for executive.
list	A data structure which allows for dynamic addition and removal of entries. It is not statically limited to a particular size.
little endian	A data representation scheme in which the bytes composing a numeric value are arranged such that the least significant byte is at the lowest address.
local	An object which was created without the GLOBAL attribute set and is accessible only on the node it was created and resides upon. In a single processor configuration, all objects are local.

local operation	The manipulation of an object which resides on the same node as the calling task.
logical address	An address used by an application. In a system without memory management, logical addresses will equal physical addresses.
loosely-coupled	A multiprocessor configuration where shared memory is not used for communication.
major number	The index of a device driver in the Driver Address Table .
manager	A group of related RTEMS' directives which provide access and control over resources.
memory pool	Used interchangeably with heap.
message	A sixteen byte entity used to communicate between tasks. Messages are sent to message queues and stored in message buffers.
message buffer	A block of memory used to store messages.
message queue	An RTEMS object used to synchronize and communicate between tasks by transporting messages between sending and receiving tasks.
Message Queue Control Block	A data structure associated with each message queue used by RTEMS to manage that message queue.
minor number	A numeric value passed to a device driver, the exact usage of which is driver dependent.
mode	An entry in a task's control block that is used to determine if the task allows preemption, timeslicing, processing of signals, and the processor mode used by the task. If the mode specifies that the task is in supervisor mode, then an interrupt level is used.
MPCI	An acronym for Multiprocessor Communications Interface Layer.
multiprocessing	The simultaneous execution of two or more processes by a multiple processor computer system.

multiprocessor	A computer with multiple CPUs available for executing applications.
Multiprocessor Communications Interface Layer	A set of user-provided routines which enable the nodes in a multiprocessor system to communicate with one another.
Multiprocessor Configuration Table	A table which contains the information needed by RTEMS only when used in a multiprocessor configuration.
multitasking	The alternation of execution amongst a group of processes on a single CPU. A scheduling algorithm is used to determine which process executes at which time.
mutual exclusion	A term used to describe the act of preventing other tasks from accessing a resource simultaneously.
nested	A term used to describe an ASR that occurs during another ASR or an ISR that occurs during another ISR.
node	A term used to reference a processor running RTEMS in a multiprocessor system.
non-existent	The state occupied by an uncreated or deleted task.
numeric coprocessor	A component used in computer systems to enhance performance in mathematically intensive situations. It is typically viewed as a logical extension of the primary processor.
object	In this document, this term is used to refer collectively to tasks, message queues, partitions, regions, semaphores, and timers.
object-oriented	A term used to describe systems with common mechanisms for utilizing a variety of entities. Object-oriented systems shield the application from implementation details.
operating system	The software which controls all the computer's resources and provides the base upon which application programs can be written.
overhead	The portion of the CPUs processing power consumed by the operating system.

packet	A buffer which contains the messages passed between nodes in a multiprocessor system. A packet is the contents of an envelope.
partition	An RTEMS object which is used to allocate and deallocate fixed size blocks of memory from an dynamically specified area of memory.
Partition Control Block	A data structure associated with each partition used by RTEMS to manage that partition.
PC	An acronym for Program Counter.
pending	A term used to describe a task blocked waiting for an event, message, semaphore, or signal.
physical address	The actual hardware address of a resource.
poll	A mechanism used to determine if an event has occurred by periodically checking for a particular status. Typical events include arrival of data, completion of an action, and errors.
pool	A collection from which resources are allocated.
portability	A term used to describe the ease with which software can be rehosted on another computer.
posting	The act of sending an event, message, semaphore, or signal to a task.
preempt	The act of forcing a task to relinquish the processor and dispatching to another task.
priority	A mechanism used to represent the relative importance of a set of items.
Program Counter	A hardware register containing the address of the current instruction.
Program Status Register	A microprocessor register typically containing the processor execution mode, interrupt level, trace mode, and condition codes.
proxy	An RTEMS control structure used to represent, on a remote node, a task which must block as part of a remote operation.

Proxy Control Block	A data structure associated with each proxy used by RTEMS to manage that proxy.
PSR	An acronym for Program Status Register.
PTCB	An acronym for Partition Control Block.
PXCB	An acronym for Proxy Control Block.
quantum	The application defined unit of time in which the processor is allocated.
queue	Alternate term for message queue. A data structure managed using the FIFO discipline.
QCB	An acronym for Message Queue Control Block.
ready	A task occupies this state when it is available to be given control of the CPU.
real-time	A term used to describe systems which are characterized by requiring deterministic response times to external stimuli. The external stimuli require that the response occur at a precise time or the response is incorrect.
reentrant	A term used to describe routines which do not modify themselves or global variables.
region	An RTEMS object which is used to allocate and deallocate variable size blocks of memory from a dynamically specified area of memory.
Region Control Block	A data structure associated with each region used by RTEMS to manage that region.
registers	Registers are locations physically located within a component, typically used for device control or general purpose storage.
remote	Any object that does not reside on the local node.
remote operation	The manipulation of an object which does not reside on the same node as the calling task.

return code	Also known as error code or return value.
resource	A hardware or software entity to which access must be controlled.
resume	Removing a task from the suspend state. If the task's state is ready following a call to the <code>t_resume</code> directive, then the task is available for scheduling.
return code	An value returned by RTEMS directives to indicate the completion status of the directive.
RNCB	An acronym for Region Control Block.
round-robin	A task scheduling discipline in which tasks of equal priority are executed in the order in which they are made ready.
RS-232	A standard for serial communications.
schedule	The process of choosing which task should next enter the executing state.
segments	Variable sized memory blocks allocated from a region.
semaphore	An RTEMS object which is used to synchronize tasks and provide mutually exclusive access to resources.
Semaphore Control Block	A data structure associated with each semaphore used by RTEMS to manage that semaphore.
shared memory	Memory which is accessible by multiple nodes in a multiprocessor system.
signal	An RTEMS provided mechanism to communicate asynchronously with a task. Upon reception of a signal, the ASR of the receiving task will be invoked.
signal set	A thirty-two bit entity which is used to represent a task's collection of pending signals and the signals sent to a task.
SMCB	An acronym for Semaphore Control Block.
SR	An acronym for Status Register.

stack	A data structure that is managed using a Last In First Out (LIFO) discipline. Each task has a stack associated with it which is used to store return information and local variables.
status code	Also known as error code or return value.
status register	A microprocessor register typically containing the processor execution mode, interrupt level, trace mode, and condition codes.
suspend	A term used to describe a task that is not competing for the CPU because it has had a <code>t_suspend()</code> directive.
synchronous	Related in order or timing to other occurrences in the system.
system call	In this document, this is used as an alternate term for directive.
target	The system on which the application will ultimately execute.
task	A logically complete thread of execution. The CPU is allocated among the ready tasks.
Task Control Block	A data structure associated with each task used by RTEMS to manage that task.
task switch	Alternate terminology for context switch. Taking control of the processor from one task and given to another.
TCB	An acronym for Task Control Block.
tick	The basic unit of time used by RTEMS. It is a user-configurable number of milliseconds. The current tick expires when the <code>tm_tick</code> directive is invoked.
tightly-coupled	A multiprocessor configuration system which communicates via shared memory.
timeout	An argument provided to a number of directives which determines the maximum length of time an application task is willing to wait for the directive to complete.

timer	An RTEMS object used to send events to the calling tasks at a later time.
Timer Control Block	A data structure associated with each timer used by RTEMS to manage that timer.
timeslicing	A task scheduling discipline in which tasks of equal priority are executed for a specific period of time before being preempted by another task.
timeslice	The application defined unit of time in which the processor is allocated.
TMCB	An acronym for Timer Control Block.
user extensions	Software routines provided by the application to enhance the functionality of RTEMS.
User Extension Table	A table which contains the entry points for each user extensions.
User Initialization Task Table	A table which contains the information needed to create and start each of the user initialization tasks.
user-provided	Alternate term for user-supplied. This term is used to designate any software routines which must be written by the application designer.
user-supplied	Alternate term for user-provided. This term is used to designate any software routines which must be written by the application designer.
vector	Memory pointers used by the processor to fetch the address of routines which will handle various exceptions and interrupts.
wait queue	The list of tasks blocked pending the release of a particular resource. Message queues, regions, and semaphores have a wait queue associated with them.
yield	When a task voluntarily releases control of the processor.

INITIAL DISTRIBUTION

	<u>Copies</u>
U.S. Army Materiel System Analysis Activity ATTN: AMXSU-MP (Herbert Cohen) Aberdeen Proving Ground, MD 21005	1
IIT Research Institute ATTN: GACIAC 10 W. 35th Street Chicago, IL 60616	1
WL/MNAG ATTN: Chris Anderson Eglin AFB, FL 32542-5434	1
Naval Weapons Center Missile Software Technology Office Code 3901C, ATTN: Mr. Carl W. Hall China Lake, CA 93555-6001	1
On-line Applications Research 3315 Memorial Parkway, SW Huntsville, AL 35801	3
CEA Incorporated Blue Hills Office Park 150 Royall Street Suite 260, ATTN: Mr. John Shockro Canton, MA 01021	1
VITA 10229 N. Scottsdale Rd Suite B, ATTN: Mr. Ray Alderman Scottsdale, AZ 85253	1
Westinghouse Electric Corp. P.O. Box 746 - MS432 ATTN: Mr. Eli Solomon Baltimore, MD 21203	1
Dept. of Computer Science B-173 Florida State University ATTN: Dr. Ted Baker Tallahassee, FL 32306-4019	1
DSD Laboratories 75 Union Avenue ATTN: Mr. Roger Whitehead Studbury, MA 01776	1

	<u>Copies</u>
AMSMI-RD	1
AMSMI-RD-GS, Dr. Paul Jacobs	1
AMSMI-RD-GC-S, Gerald E. Scheiman	1
Wanda M. Hughes	5
Phillip Acuff	4
AMSMI-RD-BA	1
AMSMI-RD-BA-C3, Bob Christian	1
AMSMI-RD-SS	1
AMSMI-RD-CS-R	15
AMSMI-RD-CS-T	1
AMSMI-GC-IP, Mr. Fred M. Bush	1
CSSD-CR-S, Mr. Frank Poslajko	1
SFAE-FS-ML-TM, Mr. Frank Gregory	1
SFAE-AD-ATA-SE, Mr. Julian Cothran	1
Mr. John Carter	1

DIST-2